

Regularization and Bayesian Learning in Dynamical Systems: Past, Present and Future

A. Chiuso*

* *Dept. of Information Engineering, University of Padova (e-mail: {chiuso}@dei.unipd.it)*

Abstract: Regularization and Bayesian methods for system identification have been repopularized in the recent years, and proved to be competitive w.r.t. classical parametric approaches. In this paper we shall make an attempt to illustrate how the use of regularization in system identification has evolved over the years, starting from the early contributions both in the Automatic Control as well as Econometrics and Statistics literature. In particular we shall discuss some fundamental issues such as compound estimation problems and exchangeability which play an important role in regularization and Bayesian approaches, as also illustrated in early publications in Statistics. The historical and foundational issues will be given more emphasis (and space), at the expense of the more recent developments which are only briefly discussed. The main reason for such a choice is that, while the recent literature is readily available, and surveys have already been published on the subject, in the author's opinion a clear link with past work had not been completely clarified.

Keywords: Identification, Learning, Numerical Methods, Linear Systems

1. INTRODUCTION

About sixty years have passed since the seminal paper by Zadeh (1956), which has coined the name “identification” and triggered research of the Automatic Control community in the broad area of data based dynamical modeling.

In the subsequent ten years the field had achieved such an importance that *IFAC* decided to start the series of *Symposiums on System Identification* (formerly *Symp. on Identification in Automatic Control Systems*) in 1967, two years after the work by Åström and Bohlin (1965) which has laid the foundations of Maximum Likelihood methods (and thus Prediction Error Methods in the Gaussian case) for ARMAX models. I have had the honor and privilege of being a plenary speaker at the 17th Symposium of the series in Beijing, and this paper has been written as a companion to the plenary presentation, which of course could not enter into many of the details that can be found here.

Despite this long history, the field is still lively and active. We believe there are two main reasons why this is so: The first is definitely the increasing importance that data centric methods are playing in many areas of Engineering and Applied Sciences with new challenges arising from the need to process high dimensional data, possibly in real time, and with little (if any) human supervision. A very nice overview of such challenges has been discussed by Lennart Ljung in his plenary at the joint IEEE Conference on Decision and Control and European Control Conference

in 2011 (Ljung et al. (2011)) and by Mario Sznajder in his SYSID 2012 semi plenary lecture, Sznajder (2012).

The second reason, which is the topic of this paper, has to do with the revitalization of (old) techniques which are rooted in the theory of regularization and Bayesian statistics.

This paper shall be focused on the role these latter techniques play in the recent developments of (linear) system identification, with the main objective to guide the reader from the early developments to the present days, with a (brief) outlook into the future. For reasons of space we will only address linear system identification, even though we believe the methods and tools discussed here have high potential in the nonlinear scenario as well (see Johansen (1997); Suykens (2001); Pillonetto et al. (2011b) and references therein). We also warn the reader that only the discrete time case will be presented. It is worth stressing that most of the recent results, just briefly discussed in Sections 4 and 7, can be framed in a continuous time scenario (see e.g. Pillonetto and De Nicolao (2010)) so that also non-uniform sampling can be handled (see e.g. Neve et al. (2008) for applications with pharmacokinetic data).

More specifically, after having introduced the problem and defined notation in Section 2, we shall provide in Section 3 an overview of parametric Maximum Likelihood/Prediction Error Methods (PEM) as formulated in Åström and Bohlin (1965). We warn the reader that *this is not* a paper about ML/PEM and thus no attempt is made to discuss its developments over the years. The main goal of Section 3 will be just to set the notation and to pinpoint the weaknesses of the parametric approach. Section 4 will

* This work has been partially supported by the FIRB project “Learning meets time” (RBFR12M3AC) funded by MIUR.

introduce the regularization approach, with an attempt to provide a complete historical overview, including early work in Statistics and Econometrics where these type of approaches have been first advocated. In order to understand the basic ideas and motivations for the Bayesian approach, Section 5 introduces the related problem of compound estimation and recalls the notion of exchangeability, as a prerequisite to Section 6 where their role in the system identification problem will be discussed. In particular the theory of compound estimation provides, from a classical (read frequentist) perspective, a sound theoretical motivation for adopting a regularization/Bayesian point of view. Finally an overview of some recent research in which I have been personally involved, regarding the design of priors and their use in structure selection problems, will be given in Sections 7 and 8 respectively. A brief outlook into the future will be provided in Section 9.

Of course this overview reflects the author's view and other people would have certainly provided a different one. Despite the long list of references I certainly have omitted many relevant ones; yet I still hope this paper can provide a starting point for anybody interested in digging a bit deeper into the roots of Bayesian Learning in System Identification.

2. STATEMENT OF THE PROBLEM

Let $u(t) \in \mathbb{R}^m$, $y(t) \in \mathbb{R}^p$ be, respectively, the measured *input* and *output* signals in a dynamical system; the purpose of system identification is to find, from a finite collection of input-output data $\{u(t), y(t)\}_{t \in [1, N]}$, a “good” dynamical model which describes the phenomenon under observation. The candidate model will be searched for within a so-called “model set” denoted by \mathcal{M} . In this paper we shall use the symbol $\mathcal{M}_n(\theta)$ for parametric model classes where the subscript n denotes the model complexity, i.e. the number of free parameters.

In this paper we shall be concerned with identification of linear models for jointly stationary processes $\{y(t), u(t)\}_{t \in \mathbb{Z}}$, i.e. models described by a convolution

$$y(t) = \sum_{k=1}^{\infty} g_k u(t-k) + \sum_{k=0}^{\infty} h_k e(t-k) \quad t \in \mathbb{Z}. \quad (1)$$

where g and h are the so-called impulse responses of the system and $\{e_t\}_{t \in \mathbb{Z}}$ is a zero mean white noise process which under suitable assumptions is the one-step-ahead prediction error; a convenient description of the linear system (1) is given in terms of the transfer functions

$$G(q) := \sum_{k=1}^{\infty} g_k q^{-k} \quad H(q) := \sum_{k=0}^{\infty} h_k q^{-k}$$

The linear model (1) yields an “optimal” (in the mean square sense) output predictor which shall be denoted later on by $\hat{y}(t|t-1)$. As mentioned above, under suitable assumptions, the noise $e(t)$ in (1) is the so-called *innovation* process $e(t) = y(t) - \hat{y}(t|t-1)$.

In order to simplify the exposition, in this work we shall only deal with feedback free (i.e. assuming that there is no feedback from y to u , see Granger (1963)) Output Error (OE) systems; thus $H(q) = I_p$ will be postulated. All ideas can be extended to handling, without major difficulties (see e.g. Pillonetto et al. (2011a)), more general situation

involving colored noise (i.e. $H(q) \neq I_p$) as well as the case where feedback is present. This however would obscure the presentation and is thus omitted.

Therefore our focus will be on linear models of the form

$$y(t) = \sum_{k=1}^{\infty} g_k u(t-k) + e(t) = \hat{y}(t|t-1) + e(t) \quad (2)$$

where (second order) joint stationarity of $\{y(t), u(t)\}_{t \in \mathbb{Z}}$ implies that $G(q)$ has to be BIBO stable (i.e. analytic outside the open unit disc of the complex plane, $|q| \geq 1$). In turn BIBO stability requires that g_k decays to zero as $k \rightarrow \infty$, and therefore the infinite summation in (2) can be approximated by a finite summation

$$y(t) \simeq \sum_{k=1}^T g_k u(t-k) + e(t) \quad (3)$$

for a large enough integer T . Since this is always possible up to an arbitrarily small approximation error¹, in this paper we shall always work with FIR models, assuming exact equality is satisfied in (3). This transforms the infinite dimensional model (2) into a finite dimensional one (3). All the results in this paper could indeed be formulated with reference to the infinite dimensional model (2), at the price of bringing so called Reproducing Kernel Hilbert Spaces (RKHS) Aronszajn (1950); Saitoh (1988) into the picture. In our opinion this only entails additional difficulties for the reader and essentially no gain in terms of tools and results and will thus be avoided. The interested reader is referred to Pillonetto and De Nicolao (2010); Pillonetto et al. (2011a, 2014). We would like to remind the reader that a detailed study of the asymptotic properties of Bayes procedures for infinite dimensional models is delicate and outside the scope of this paper, see for instance Knapik et al. (2011).

In the following we shall use the notation $Y \in \mathbb{R}^{pN}$ to denote the (column) vector containing the stacked outputs $y(t)$, $t \in [1, N]$ and $\hat{Y}(g)$ the vector of stacked predictors $\hat{y}(t|t-1)$, $t \in [1, N]$ which is a linear function of the impulse response coefficients g_k :

$$\hat{Y}(g) = \Phi g$$

where $\Phi \in \mathbb{R}^{pN \times pmT}$ is a suitable matrix built with the input data $u(t)$, while the column vector $g \in \mathbb{R}^{pmT}$ contains the (vectorized) impulse response matrix coefficient $g_k \in \mathbb{R}^{p \times m}$, $k \in [1, T]$.

Performing identification of g (i.e. estimation of the impulse response from a finite set of input output data) can be thus framed as estimation of the unknown g in the linear model

$$Y = \Phi g + E \quad g \in \mathbb{R}^d \quad d := pmT \quad (4)$$

Unfortunately the dimension d of the unknown vector g may be very large (and possibly much larger than the length of the available data Y) so that the inverse problem

¹ Rigorously one should account for the transient effect, which can be beneficial for small data sets where N (and thus implicitly T) is necessarily small. This can be done rather easily estimating the free response for each output channel. Details are outside the scope of this paper and shall not be discussed here.

of determining g from Y in (4), e.g. minimizing the square loss

$$\hat{g}_{LS} := \arg \min_g \|Y - \Phi g\|^2, \quad (5)$$

may be (very) ill conditioned. This may be due to the fact that the input process lives in a high dimensional space (pm large, so that many impulse responses need to be estimated) or simply because g_k decays very slowly to zero and thus many lags need to be included in the parameter vector g .

To face this problem one has to impose constraints on the structure of the vector g . One possibility is to parametrize $g_k = g_k(\theta)$ using a vector $\theta \in \mathbb{R}^n$, $n \ll pN$. In the remaining part of the paper we shall sometimes make explicit the dimension of the parameter vector using the notation θ_n . For instance one may assume $G(q, \theta) := \mathcal{Z}[g_k(\theta)](q)$ is a rational function of a given McMillan degree or $G(q, \theta)$ is expressed via an orthonormal basis expansion (Wahlberg (1991)). When parametric models are considered the notation $\hat{y}_\theta(t|t-1)$ will be used for the predictor. We shall come back to these representations later in the paper.

Note however that even “parametric” representations may require a large number (possibly infinite) of parameters, such as orthogonal basis expansions. For instance if the Fourier basis q^{-k} , $k = 1, \dots, m$ is chosen, a finite expansion has the form $G(q, \theta) = \sum_{k=1}^m \theta_k q^{-k} = \sum_{k=1}^m g_k q^{-k}$. Apparently the “parameter” $\theta = g \in \mathbb{R}^{mpT}$ needs to grow in size in order to approximate any transfer matrix G . We shall refer to this latter cases as “non-parametric” (even though, strictly speaking, $g = \theta$ is still a parameter) since the length of the parameter vector $g = \theta$ may need to be arbitrarily large irrespectively on the number of data N to guarantee that the model class $G(q, \theta)$ accurately describes the “true” underlying system.

Whether g is described in a parametric fashion or in a non parametric one, the model “complexity” (the complexity of the model class \mathcal{M} where the estimator \hat{g} is searched for) needs to be controlled so that the estimated model behaves “well” on future data (e.g. in terms of output simulation). Therefore, given Y_{test} new data *not* used for identification, it is to be chosen so that the so-called *Mean Squared Error*

$$MSE_Y(\hat{g}) = \lim_{N \rightarrow \infty} \mathbb{E} \frac{1}{N} \|Y_{test} - \hat{Y}(\hat{g})\|^2 \quad (6)$$

is small. In the particular case of Output Error models with white (unit variance) noise input, the Mean Squared Error (7) reduces to

$$MSE_Y(\hat{g}) = \mathbb{E} \|g - \hat{g}\|^2 + \sigma^2. \quad (7)$$

This is nothing but the usual *bias-variance tradeoff*: the model has to be rich enough to capture the “true” data generating mechanism (low bias) but also simple enough to be estimated using the available data with small variability (low variance). The loss in (7) is called a *compound* loss on the (possibly infinite) vector g . As we shall see later on this plays a fundamental role in studying the properties of regularized estimators.

3. PARAMETRIC METHODS: THE MAXIMUM LIKELIHOOD/PEM ERA

The nuts and bolts of Prediction Error Methods (PEM) (or Maximum Likelihood when the innovations $e(t) = y(t) -$

$\hat{y}(t|t-1)$ are Gaussian) were laid down in the System Identification community by the seminal paper Åström and Bohlin (1965) considering SISO ARMAX models. In the simplified setup we consider in this paper this corresponds to $m = p = 1$, and the model class $\mathcal{M}_n(\theta_n)$ being described in terms of rational transfer functions $G(q, \theta_n)$ of degree ν :

$$G(q, \theta_n) = \frac{\sum_{i=0}^{\nu} b_i q^i}{q^{\nu} + \sum_{i=0}^{\nu-1} a_i q^i} \quad \sigma^2 = \text{Var}\{e(t)\}$$

$$\theta_n := [a_0, \dots, a_{\nu-1}, b_0, \dots, b_{\nu}, \sigma^2]^T \in \mathbb{R}^n \quad n = 2\nu + 2. \quad (8)$$

The estimator $\hat{\theta}_n$ was found in Åström and Bohlin (1965) following the maximum likelihood approach, which under the assumption of Gaussian innovation is equivalent to solving a (non linear) least squares problem. Let $p_{\theta_n}(Y)$ be the likelihood function of the data Y under the n -dimensional parameter vector θ_n , we have:

$$\begin{aligned} \hat{\theta}_n &:= \arg \max_{\theta \in \mathbb{R}^n} p_{\theta}(Y) \\ &= \arg \min_{\theta \in \mathbb{R}^n} -2 \log(p_{\theta}(Y)) \end{aligned} \quad (9)$$

In Åström and Bohlin (1965) also attention was paid to numerical procedures for performing likelihood maximization, which exploit the structure of the underlying model. The idea of using Maximum Likelihood for estimating time series models can be traced back to Wald and Mann (1943), see also Cramér (1946); Grenander (1950); Whittle (1953). The theory of Maximum Likelihood estimation allows to derive asymptotic results such as consistency and asymptotic statistical efficiency under correct specification of the model class $M_n(\theta)$. Åström and Bohlin (1965) also discuss testing procedures for checking the choice of model order, based on (estimates of) the Hessian of the log likelihood $L(\theta_n) := -2 \log(p_{\theta_n}(Y))$ and on testing whiteness of the residual sequence $e_{\hat{\theta}_n}(t) := y(t) - \hat{y}_{\hat{\theta}_n}(t|t-1)$.

The literature on Statistical System Identification based on the Maximum Likelihood/Prediction Error framework has since then been extensively developed in the Statistics, Econometrics and Control literatures and it is fair to say that it has reached by now a reasonable maturity, as testified by the highly cited textbooks Ljung (1999); Söderström and Stoica (1989); Box et al. (2008); Brockwell and Davis (1991); Hannan and Deistler (1988). It is worth here to recall that much (but possibly not enough) work has been devoted to the problem of selecting the model class (here the “model complexity” n). The early literature has focused on testing approaches, see Quenouille (1947), Wold (1949) and Åström and Bohlin (1965). Penalty based criteria have been first advocated in Akaike (1974). The general form is

$$\hat{n} := \arg \min_n L(\hat{\theta}_n) + c(N)n \quad (10)$$

where $c(N)$ is a (non decreasing) function of the data length. Different choices lead to different criteria: $c(N) = 2$ is known as AIC (Akaike (1974)), $c(N) = \log(N)$ leads to BIC developed in Schwarz (1978) and Rissanen (1978); other alternatives are possible, see e.g. Hurvich and Tsai (1989). The properties of these order estimation criteria and of the estimators $\hat{\theta}_{\hat{n}}$ obtained after model selection

has been performed, called also *Post Model Selection Estimators* (PMSE), have been studied by a number of authors. We shall here mention only two very relevant contributions: Yang (2005) discusses the relative merits of AIC and BIC type procedures, i.e. consistency and optimal minimax rate of convergence and their mutual exclusiveness; Leeb and Pötscher (2005, 2006) show also that the asymptotic properties of PMSE are non-uniform w.r.t the “true” model which generates the data, thus making the finite sample properties of the PMSE $\hat{\theta}_n$ possibly very different to the predictions of asymptotic theory. It is also worth mentioning that, for finite data size, the “true” order is not necessarily the “best” one (e.g. in terms of facing the bias-variance tradeoff).

We conclude this section by observing that the model selection step is, in our opinion, a rather critical aspect of the identification procedure. This is in fact in line with the findings, amongst others, of Leeb and Pötscher (2005, 2006) and with the experimental observations in Pillonetto and De Nicolao (2010); Pillonetto et al. (2011a); Chen et al. (2012).

4. NONPARAMETRIC METHODS, REGULARIZATION AND BAYESIAN APPROACHES: FORMULATION AND HISTORICAL OVERVIEW

One of the main drawbacks of parametric approaches, as discussed above, is the need to estimate the model structure (and its complexity n) prior to computing a point estimate of the parameter θ_n . On the other hand working in a nonparametric scenario there is no need to perform “parametric” model selection; the truncation length T of the impulse response or the number L of coefficients in a basis expansion are just supposed to be “large” so that the “true” model is close enough to the model class. However one needs to find alternative ways to control the model complexity such as to make the inverse problem of determining g from Y in (4) well-posed.

Thus a different line of work has studied system identification in the context of ill-posed (non-parametric) linear regression models of the form (4); the main idea can be described either in a probabilistic framework (Bayesian formulation) or deterministic one (regularization).

From the regularization perspective one can consider estimators of the form

$$\hat{g}_\eta^{Reg} = \arg \min_g \|Y - \Phi g\|^2 + J_\eta(g) \quad (11)$$

for some “penalty” function $J_\eta(g) \geq 0$, possibly depending on some parameters η , which attempts to “discourage” certain undesired solutions. Alternatively, the same estimator \hat{g}_η can be obtained in a Bayesian framework assuming that $E := Y - \Phi g$ is a zero mean Gaussian vector with variance $\sigma^2 I$ and g has a prior

$$p_\eta(g) \propto e^{-\frac{1}{2} J_\eta(g)} \quad (12)$$

so that the Maximum a Posteriori (MAP) estimator

$$\hat{g}_\eta^{MAP} = \arg \max_g p_\eta(g|Y) = \arg \max_g p(Y|g)p_\eta(g), \quad (13)$$

equals \hat{g}_η^{Reg} in (11). We shall thus use the notation $\hat{g}_\eta := \hat{g}_\eta^{MAP} = \hat{g}_\eta^{Reg}$ hereafter.

As we shall see an interesting case is when $J_\eta(g)$ is a quadratic form

$$J_\eta(g) = g^\top K_\eta^{-1} g \quad K_\eta = K_\eta^\top > 0 \quad (14)$$

where now the matrix K_η is possibly parametrized by a vector of parameters (called hyperparameters, see later on) η . Under this assumption the density $p_\eta(g)$ in (12) is Gaussian with zero mean and variance K_η and the MAP estimator (13) equals the posterior mean

$$\hat{g}_\eta := \mathbb{E}_\eta[g|Y] \quad (15)$$

which is to be computed for fixed η . The matrix K_η is also known in the Machine Learning literature as “kernel”, see Rasmussen and Williams (2006); Schölkopf and Smola (2001). Under the assumption (14) the regularized inverse problem (11) is known as Tikhonov-regularization (see Tikhonov and Arsenin (1977); Melkman and Micchelli (1979)) in the literature of inverse problems (see also Bertero (1989)) and as ridge regression in the statistics and numerical analysis literature Riley (1955); Foster (1961); Hoerl (1962); Hoerl and Kennard (1970); Marquardt and Snee (1975)

$$\hat{g}_\eta = \arg \min_g \|Y - \Phi g\|^2 + g^\top K_\eta^{-1} g \quad (16)$$

The model class is here described as the set of impulse response g such that the term $g^\top K_\eta^{-1} g$ is “small”. The size of this set is related to the singular values profile of K_η , which are also related to the concepts of *degrees of freedom*, see Hastie et al. (2009) and equation (3.12) in Pillonetto and Chiuso (2015b).

Note that the *shape* as well as the *size* of this set depends on the “regularization (hyper)parameter” η , which needs to be fixed in order to compute the estimator (16). This is a way to control the model complexity (the “size” of the model class) which can be done trading *bias* and *variance* controlling the *generalization error* (7). However, since the Mean Squared Error cannot be computed unless the true system is known, estimators have to be found. This can be done by so called *Cross-Validation*, i.e. estimating (7) using data directly (see e.g. Hastie et al. (2009)), or by estimators of the MSE such as *Stein Unbiased Risk Estimators* (SURE) Stein (1981); also methods based on “deterministic/worst case” arguments have been considered in the context of identification, see e.g. Tempo (1995).

One nice feature of the Bayesian formulation (13) as opposed to the regularization one (11) is that the former provides tools to estimate the hyperparameters η using the so-called marginal likelihood $p_\eta(Y)$, i.e. the likelihood of the hyper parameters η once the unmeasurable quantities (g) have been integrated out. Let us denote with

$$\hat{\eta} := \arg \max_\eta p_\eta(Y) \quad (17)$$

the marginal likelihood estimator of η . Then plugging (17) in (15) we obtain the so-called empirical Bayes estimator (see Robbins (1956, 1964); Good (1965); Maritz and Lwin (1989); Rasmussen and Williams (2006); Aravkin et al. (2014); Pillonetto and Chiuso (2015a)) of g :

$$\hat{g}_{EB} := \mathbb{E}_{\hat{\eta}}[g|Y] \quad (18)$$

which is of course a function of the chosen prior $p_\eta(g)$ and thus of the kernel K_η .

To the best of our knowledge Bayesian methods (13) for estimating dynamical systems have first been advocated in Leamer (1972) and Shiller (1973) where an FIR model

(called *distributed lag* in those references) of the form (3) was considered. In fact Leamer (1972) and Shiller (1973) were the first to talk about (and apply) “Bayesian” methods for system identification, arguing that “rigid parametric” structures may be inadequate to the purpose; the reader may consult instead Tiao and Zellner (1964) for an early overview on the use of Bayes priors in the context of linear regression. These ideas can be traced back, to the best of the author’s knowledge, to an early paper by Whittaker (1922), where smoothing techniques (actually Bayesian priors) were first advocated for the purpose of “denoising” some measured data, with particular reference to the field of Actuarial Science in this early reference.

In Shiller (1973) it is assumed that g is a normal random variable with zero mean and covariance matrix K , so that the posterior mean of g (conditioned on the observations Y) is

$$\begin{aligned}\hat{g} &= \mathbb{E}[g|Y] = \text{cov}(g, Y) [\text{Var}(Y)]^{-1} Y \\ &= K (\Phi K \Phi^\top + \sigma^2 I)^{-1} Y = (\Phi^\top \Phi + \sigma^2 K^{-1})^{-1} \Phi^\top Y;\end{aligned}\quad (19)$$

The latter (see the rightmost equation) has the form of a (generalized) ridge regression estimator, obtained from (16) with $K_\eta = \frac{1}{\sigma^2} K$. This is a well known technique to “stabilize” ill-conditioned linear inverse problems (see Riley (1955); Foster (1961); Hoerl (1962); Hoerl and Kennard (1970); Marquardt and Snee (1975)).

It is clear from equation (19) that the prior covariance K plays the role of the “ridge” matrix in (16), showing the equivalence between Bayesian estimation in the Gaussian framework and Ridge Regression.

The ridge matrix K was chosen in Shiller (1973) so as to guarantee some degree of smoothness of g . However no prior information was enforced on the fact that g is expected to be an (exponentially) decaying sequence.

An important step in this direction was then made in Akaike (1979) who criticized the choice made by Shiller (1973) of describing the unknown impulse response imposing a smoothness prior. Instead Akaike (1979) suggests that smoothness should be enforced on the frequency response

$$G(e^{j\omega}) := \sum_{k=1}^T g_k e^{j\omega k}$$

for instance via the integral of the (squared) first derivative $\frac{dG(e^{j\omega})}{d\omega}$

$$\frac{1}{2\pi} \int_0^{2\pi} \left| \frac{dG(e^{j\omega})}{d\omega} \right|^2 d\omega = \sum_{k=1}^T k^2 |g_k|^2.$$

Using this as a penalty term leads to the regularized inverse problem:

$$\hat{g} := \arg \min_g \|Y - \Phi g\|^2 + \gamma \sum_{k=1}^T k^2 |g_k|^2, \quad (20)$$

again in the form of a ridge regression estimator. As discussed above this is completely equivalent to assuming a zero mean Gaussian prior on g as in Shiller (1973), yet with a decaying diagonal covariance matrix

$$K_\gamma := \frac{1}{\gamma} \text{diag} \left\{ 1, \frac{1}{4}, \frac{1}{9}, \dots, \frac{1}{T^2} \right\} \quad (21)$$

It is also worth observing that Akaike (1979) is the first to advocate the use of the Empirical Bayes approach (probably first introduced in Robbins (1956, 1964) and Good (1965) under the name of *Type II maximum Likelihood*) to estimate the scaling factor γ . More specifically γ is estimated as²:

$$\hat{\gamma} := \arg \max_{\gamma} \int p(Y|g)p(g|\gamma) dg \quad (22)$$

and then substituted in (20). Note that with Gaussian errors E (so that $p(Y|g)$ is a Gaussian with mean Φg and variance $\sigma^2 I$) and Gaussian prior $p(g|\gamma)$ the integral in (22) is available in closed form. It is interesting to observe that Akaike suggested to use the marginal likelihood also to estimate other system properties not captured by the covariance (21) such as possible delays.

This very same approach was then followed by a series of works by Kitagawa and Gersh, see e.g. Kitagawa and Gersh (1984, 1985) which are well documented in the book Kitagawa and Gersch (1996).

A related formulation, yet with a slightly different goal, is found in the so called “stochastic embedding” approach by Goodwin and Salgado (1989); Goodwin et al. (1992), see Ljung et al. (2014) for details on a “modern” interpretation. Essentially the idea was to admit that the “true” transfer function $G(q)$ is only partially captured by the chosen model class $\mathcal{M}_n(\theta)$ so that

$$G(q) = G(q, \theta_0) + \tilde{G}(q) \quad G(q, \theta_0) \in \mathcal{M}_n(\theta) \quad (23)$$

and $\tilde{G}(q)$ represents a model error. Goodwin et al. (1992) describe this model error using a “stochastic” approach similar to Akaike (1979), yet with an exponentially decaying covariance function. More specifically:

$$\tilde{G}(q) := \sum_{k=1}^T \tilde{g}_k q^{-k}$$

where $\tilde{g} := [\text{vec}(\tilde{g}_1)^\top, \dots, \text{vec}(\tilde{g}_T)^\top]^\top \sim \mathcal{N}(0, K)$ with

$$K_{\rho, \lambda} := \lambda \text{diag} \{1, \rho, \rho^2, \dots, \rho^{T-1}\} \quad (24)$$

The model error $\tilde{G}(q)$ is estimated in Goodwin et al. (1992) starting from the least squares residuals $v_{\hat{\theta}}(t) := y(t) - G(q, \hat{\theta})u(t)$ where

$$\hat{\theta} = \arg \min_{\theta} \sum_{k=1}^T \|y(t) - G(q, \theta)u(t)\|^2$$

Under the assumption (23), the error $v_{\hat{\theta}}(t) := y(t) - G(q, \hat{\theta})u(t)$ is expected to be described by the model

$$v(t) = \tilde{G}(q)u(t) + e(t)$$

and thus Goodwin et al. (1992) propose to estimate λ and ρ from marginal likelihood maximization

$$(\hat{\lambda}, \hat{\rho}) := \arg \max_{\lambda, \rho} \int p(V_{\hat{\theta}}|\tilde{g})p(\tilde{g}|\lambda, \rho) d\tilde{g} \quad (25)$$

where $V_{\hat{\theta}} := [v_{\hat{\theta}}^\top(1), \dots, v_{\hat{\theta}}^\top(N)]^\top$. This leads to an empirical Bayes estimator similar to that utilized in Akaike (1979). Note however that Goodwin et al. (1992) is not focused on estimating the transfer function $\tilde{G}(q)$ but rather its variance, which is utilized to describe statistically the

² Slight variations suggested by Akaike include marginalization with respect to the noise variance σ^2 , subject to improper priors.

model error. This in turn is used to introduce an order estimation criteria based on minimization of an estimate of the Mean Squared Error

$$\mathbb{E}\|G(q) - G(q, \hat{\theta}_n)\|^2$$

while accounting for the model error model $\tilde{G}(q)$ (see e.g. equations (89-92) in Goodwin et al. (1992)). It is worth to observe that the stochastic embedding approach, which postulates a white noise model with exponentially decaying variances (24) on the impulse response sequence, implies a Lipschitz condition on the impulse response:

$$\begin{aligned} \mathbb{E}[\|G_\lambda(e^{j\omega_1}) - G_\lambda(e^{j\omega_2})\|^2] &\leq f(\lambda)(\omega_1 - \omega_2)^2 \\ \lim_{|\lambda| \rightarrow 1^-} f(\lambda) &= +\infty \end{aligned} \quad (26)$$

similar in spirit to the “frequency domain” smoothness condition advocated in Akaike (1979).

A significant amount of work with Bayesian flavor is also found in the econometrics literature. For instance in a series of works Litterman, Doan and Sims (see e.g. Doan et al. (1984) for an overview) study time varying multivariate autoregressive models of the form³

$$y(t) = \sum_{k=1}^m A_{k,t} y(t-k) + C(t) + \epsilon(t) \quad (27)$$

where $C(t)$ is a time varying deterministic sequence and $\epsilon(t)$ is a white noise sequence.

The coefficients $A_{k,t}$ and $C(t)$ are modeled as a function of a number of hyperparameters (called π_i , see page 12 of Doan et al. (1984)) as follows: The time evolution is modeled as a first order auto regression

$$\text{vec}(A_{k,t+1}) = \pi_8 \text{vec}(A_{k,t}) + (1 - \pi_8) \text{vec} \bar{\mu} + \mu(t)$$

where $\bar{\mu}$ is a fixed vector and $\mu(t)$ is a known time varying sequence. The initial conditions of this recursion (mean and variances) are fixed (see eq. (4-6) Doan et al. (1984)). In particular the coefficients $[A_{k,0}]_{i,j}$ are supposed to satisfy

$$\text{Var}\{[A_{k,0}]_{i,j}\} = \begin{cases} \frac{\pi_5 \pi_1}{k e^{\pi_4 w_{ii}}} & i = j \\ \frac{\pi_5 \pi_2 \sigma_i^2}{k e^{\pi_4 w_{ij}} \sigma_j^2} & i \neq j \end{cases} \quad (28)$$

$$\text{Cov}\{[A_{k,0}]_{i,j}, [A_{h,0}]_{m,n}\} = \text{Var}\{[A_{k,0}]_{i,j}\} \delta_{hk} \delta_{im} \delta_{jn}$$

where δ_{hk} is the Kronecker delta, w_{ij} are suitable user defined weights and σ_i^2 are the variances of the prediction error for the i -th component of $y(t)$. This is one form of the so called *Minnesota prior*, which has been discussed quite extensively in the econometrics literature; several variations and extensions are found, see for instance Lütkepohl (2007); Giannone et al. (2015).

It is apparent that the resulting prior distribution has a considerable flexibility; it is reasonable to expect that the performance of any algorithm based on such prior may heavily depend upon the (rather arbitrary) user choices; Doan et al. (1984) discuss extensively on the meaning of each of these choices and give suggestion as to how these should be made; delving into these details is outside the scope of this lecture. Suffices here to notice that, assuming the coefficients are fixed over time, equations (28) imply a diagonal prior covariance matrix with elements decaying

as $\frac{1}{k}$ where k is the lag index while different versions of the Minnesota Prior suggest decay rates of the form $\frac{1}{k^\alpha}$; this clearly reminds of the prior suggested by Akaike (1979), in which $\alpha = 2$. Note that $\alpha > 1$ is required in order to guarantee that realizations g_k , $k = 1, \dots, T$ are (a.s.) in ℓ_1 as $T \rightarrow \infty$.

The econometrics literature has since then studied Bayesian procedures for system identification rather intensively, mostly under the acronym *Bayesian VARs*; the main driving motivation was that of handling high dimensional time series (i.e. p large, called *cross sectional dimension* in the econometrics literature) with possibly many explicative variables (m large), see for instance Knox et al. (2001); De Mol et al. (2008); Banbura et al. (2010); Giannone et al. (2015). While Doan et al. (1984) and Banbura et al. (2010) propose tuning the hyperparameters using out-of-sample and in-sample error respectively, Knox et al. (2001) and the most recent work Giannone et al. (2015) adopt an Empirical Bayes approach using the marginal likelihood (i.e. the likelihood function of the data Y given the hyperparameters once the unknown parameters have been integrated out) for hyperparameter estimation; Giannone et al. (2015) claims the superiority of this approach w.r.t. previous “ad-hoc” techniques Doan et al. (1984); Banbura et al. (2010).

The most recent developments in the area of system identification, which have followed the seminal contribution Pilonetto and De Nicolao (2010) and the subsequent papers Pilonetto et al. (2011a); Chen et al. (2012); Chiuso and Pilonetto (2012) (see also the survey Pilonetto et al. (2014)) are based on modeling the unknown impulse response g as a zero mean Gaussian process with a suitable covariance matrix K_η . Most of the work has been focused on designing the matrix K_η (see Pilonetto and De Nicolao (2010); Pilonetto et al. (2011a); Chiuso and Pilonetto (2012); Chen et al. (2012, 2014); Chen and Ljung (2015)) as well as on studying the properties of the Empirical Bayes estimator (18), Aravkin et al. (2014); Pilonetto and Chiuso (2015a). More details will be given later on.

5. COMPOUND DECISION AND ESTIMATION PROBLEMS AND EXCHANGEABILITY

In this section we shall attempt to bridge the gap between the classical “frequentist” view in statistics, which is rooted in the idea of “repeated” experiments with fixed but unknown parameters, and the Bayesian view where parameters are assumed to satisfy a prior distribution and are estimated minimizing suitable risk functionals.

Our detour goes back to early work by Robbins (1951, 1956, 1964), Stein (1956); James and Stein (1961), Good (1965), Lindley and Smith (1972), Efron and Morris (1973).

Two mind blowing papers in statistics have been Robbins (1951) and Stein (1956) which studied respectively the so called compound decision and estimation problems. The main message is as follows: assume there is a parameter vector $\alpha \in \mathbb{R}^B$ which we need to make a decision on or estimate. Two possibilities are that either the components of α are binary variables (e.g. $\alpha_i \in \{-1, 1\}$) (Robbins (1951, 1956)) or continuous variables $\alpha_i \in \mathbb{R}$ (Stein (1956); James and Stein (1961)). For simplicity of exposition, let

³ Exogenous inputs $u(t)$ have also been considered in this framework.

us consider the estimation problem and assume measurements are available of the form

$$Y = \alpha + E \quad (29)$$

where E is a Gaussian noise with zero mean and covariance $\sigma^2 I$ (assume for simplicity that σ^2 is known, but similar arguments hold if σ^2 needs to be estimated from data).

Assume that, given a rule $\hat{\alpha} := \delta(Y)$ used to estimate α , one measures the loss incurred in the “compound” manner:

$$L(\alpha, \delta) := \frac{1}{B} \sum_{i=1}^B (\alpha_i - \delta_i(Y))^2 \quad (30)$$

This means that we are not interested in the error $\alpha_i - \delta_i(Y)$ on each component, but only on the aggregate error (30) which is symmetric (i.e. permutation invariant). Note that in the System Identification scenario one is typically interested exactly in compound losses of this form, see (7). We now need to define the concepts of *admissibility*.

Definition 1. An estimator δ is *admissible* (w.r.t. the loss function $L(\alpha, \delta)$) if there does not exist any other estimator δ^* such that

$$\mathbb{E}L(\alpha, \delta^*) \leq \mathbb{E}L(\alpha, \delta) \quad \forall \alpha$$

with strict inequality for some α . If such an estimator δ^* exists, δ is inadmissible and δ^* is said to dominate δ .

Robbins (1951, 1956) has shown that, for the decision problem $\alpha_i \in \{-1, 1\}$, the “classical” Neyman-Pearson test (Lehmann and Romano (2005))

$$\hat{\alpha}_i = \delta_i^{NP}(Y) = \text{sign}(Y_i) \quad (31)$$

is not admissible. Similarly, Stein (1956) has shown that, for the estimation problem $\alpha_i \in \mathbb{R}$, the least squares estimator

$$\hat{\alpha}_i = \delta_i^{LS}(Y) = Y_i \quad (32)$$

is *inadmissible* (for $B > 2$). Estimation (or decision) rules which dominate $\delta_i^{NP}(Y)$ and $\delta_i^{LS}(Y)$ are provided in Robbins (1956) and James and Stein (1961) respectively. We shall now study in more detail the case considered in Stein (1956); James and Stein (1961).

The most famous example of such rules which dominate Least Squares is the so-called James-Stein estimator (discovered by James, which was at the time a student of Stein, James and Stein (1961)):

$$\hat{\alpha} = \delta^{JS}(Y) = \left(1 - \frac{(B-2)\sigma^2}{\|Y\|^2}\right) Y \quad (33)$$

It is clear that the estimation rule $\delta^{JS}(Y)$ has a “compound” flavor, in that the estimator $\delta_i^{JS}(Y)$ of the i -th component α_i depends on all the measurements Y and not only on Y_i . This may be surprising given that the noise E in (29) has uncorrelated components. It is an easy calculation to show that

$$\mathbb{E}L(\alpha, \delta^{JS}) < \mathbb{E}L(\alpha, \delta^{LS}) = \sigma^2 \quad \forall \alpha : \|\alpha\| < \infty$$

provided $B > 2$. Efron and Morris (1973) discuss the link with the empirical Bayes approach. We derive here a simplified analysis. Let us postulate a Gaussian prior on $\alpha \sim \mathcal{N}(0, \lambda I)$ and assume the noise variance σ^2 is known. The Bayes estimator (15) takes the form

$$\hat{\alpha}_\lambda := \mathbb{E}_\lambda[\alpha|Y] = \frac{\lambda}{\lambda + \sigma^2} Y \quad (34)$$

and the hyperparameter λ can be estimated using (17) leading to

$$\hat{\lambda} = \max\left(\frac{1}{B}\|Y\|^2 - \sigma^2, 0\right) \quad (35)$$

Thus leading to the Empirical Bayes estimator

$$\hat{\alpha}_{EB} := \hat{\alpha}_{\hat{\lambda}} = \left(1 - \frac{B\sigma^2}{\|Y\|^2}\right)^+ Y \quad (36)$$

where, given $v \in \mathbb{R}^B$, the positive part operator $(v)^+$ is defined as $(v)_i^+ = \max(v_i, 0)$. The estimator (36) has the same form of the positive-part version of (33) (see James and Stein (1961)) with a slightly different numerator (B in lieu of $B-2$), which can shown to dominate the Least Squares estimator (32) provided B is large enough.

Remark 2. An alternative “Empirical Bayes” derivation of the James Stein estimator (33) goes as follows: Instead of maximizing the marginal posterior distribution of Y , one can observe that, $\|Y\|^2 \sim (\sigma^2 + \lambda)\chi_B^2$, so that $\hat{A} := \sigma^2 \frac{B-2}{\|Y\|^2}$ is an unbiased estimator of $A = \frac{\sigma^2}{\sigma^2 + \lambda}$. Then, replacing \hat{A} for A in the expression $\hat{\alpha}_\lambda := \mathbb{E}_\lambda[\alpha|Y] = \frac{\lambda}{\lambda + \sigma^2} Y = (1 - A)Y$ (see (34)), gives:

$$\hat{\alpha}'_{EB} = \left(1 - \hat{A}\right) Y = \left(1 - \sigma^2 \frac{B-2}{\|Y\|^2}\right) Y = \delta^{JS}(Y)$$

The overall message we learn from this example is that ridge regression/Bayesian estimation can provide valuable tools to build estimators which dominate standard (non-regularized) Least Squares provided one is interested in compound losses (30) and not in the quality of a single parameter α_i . In addition the Empirical Bayes paradigm provides also a tool for selecting, in the class of Bayes estimators (34), a specific one which dominates Least Squares. The Empirical Bayes rule leads to the “compound” estimator (36) which couples all observations similarly to James and Stein (1961). Unfortunately the situation is rather tricky if one considers measurement models such as

$$Y = S\alpha + E \quad (37)$$

and more general (quadratic) loss functions (30) such as:

$$L(\alpha, \delta) := (\delta - \alpha)^\top Q (\delta - \alpha) \quad (38)$$

The statistics literature is rich of contributions in this direction which is certainly impossible to survey here. The reader is referred to Strawderman (1978); Berger (1982); Casella (1980); Maruyama and Strawderman (2005) and references therein. We shall come back to this issue in Section 6 where such connection will be studied in some detail.

Before concluding this section there is one more important remark to make which has to do with the concept of exchangeability, see de Finetti (1931); Hewitt and Savage (1955); Lindley and Smith (1972).

Definition 3. We say that an infinite sequence of random variables x_i , $i \in \mathbb{N}$ is *exchangeable* if $\forall h \geq 2$ the joint distribution of x_1, \dots, x_h is invariant to permutations of the indexes set $1, \dots, h$.

An important characterization of exchangeable random sequences is due to de Finetti (1931) and extended in Hewitt and Savage (1955) who have shown that an infinite sequence of random variables x_i , $i = 1, \dots$, is exchangeable if and only if there exists a random distribution $\Pi(x)$ such that the joint distribution $P(x_1, \dots, x_h)$ satisfies

$$P(x_1, \dots, x_h) = \int \prod_{i=1}^h \Pi(x_i) Q(d\pi) \quad (39)$$

for a suitable measure $Q(\pi)$. This is to say that, conditionally on the distribution Π , the x_i 's are i.i.d. with distributions $\Pi(x)$.

For the setup considered in Robbins (1951) and Stein (1956), exchangeability of α_i is a reasonable assumption which is also reflected in the loss function $L(\alpha, \delta)$ being insensitive to permutations of α_i 's. The result in equation (39) essentially says that, for an exchangeable sequence, it is reasonable to assume that there is a hierarchical prior model for the joint distribution. Further, results in de Finetti (1931); Hewitt and Savage (1955) suggest that the distribution $\Pi(x)$ can be estimated from a sample x_1, \dots, x_h . When the x_i 's are not observed, but only their (noisy) measurements y_i are available, then $\Pi(x)$ is to be estimated from y_1, \dots, y_h . The reader may consult also Edelman (1988) where further discussion concerning the estimation of Π as well as theoretical results on the resulting estimator's properties are derived.

The Empirical Bayes estimator (36) can be framed in this context, provided one postulates that $\frac{d\Pi}{dx}$ is the density of a zero mean Gaussian random variable with variance λ . Then λ can be estimated from the sample y_1, \dots, y_h as in (35). The next Section will elaborate upon this in some detail.

6. CONNECTIONS BETWEEN COMPOUND ESTIMATION PROBLEMS AND SYSTEM IDENTIFICATION

We shall now see how the system identification problem fits in the framework of compound estimation discussed in the previous Section. As anticipated at the end of Section 5, we shall need to consider measurement models of the form (37) and "weighted" losses of the form (38).

Consider now the linear model (4) and further assume that the unknown vector g is expanded w.r.t. a (possibly orthogonal) basis ψ_i , such that $\|\psi_i\|_{\ell^1} = \mu_i$, $i = 1, \dots, B$:

$$g = \sum_{i=1}^B \psi_i \alpha_i \quad (40)$$

For convenience of notation we define $\Psi := [\psi_1, \dots, \psi_B]$ and $\alpha := [\alpha_1, \dots, \alpha_B]^\top$. With such notation, and using (40), model (4) can be rewritten in the form

$$Y = \Phi \Psi \alpha + E \quad (41)$$

which is of the form (37) with $S := \Phi \Psi$.

We shall come back in Section 7 on the issue of selecting the basis vectors ψ_i . Assume also that the sequence μ_i measures the "importance" of the i -th basis vector ψ_i and is such that, in the limit⁴ as $B \rightarrow \infty$,

$$\sum_{i=1}^{\infty} |\mu_i| < \infty \quad (42)$$

so that the sum (40) converges (in ℓ_1) provided $|\alpha_i| < c$, $\forall i$ and converges in mean square if the α_i 's are i.i.d. with zero means and finite variance.

⁴ Note that here, in order to take the limit as $B \rightarrow \infty$, we should think of g as an element of $\ell(\mathbb{Z}^+)$.

Note that, if one postulates that $\alpha \sim \mathcal{N}(0, \lambda I)$, then

$$g = \Psi \alpha \sim \mathcal{N}(0, \lambda K) \quad K := \Psi \Psi^\top \quad (43)$$

where now the columns of Ψ are the (unnormalized) eigenvectors of K and (40) is the (normalized) Karhunen-Loève (Loève (1963)) expansion of g . This leads back to the standard zero mean Gaussian prior on the unknown g adopted in the Bayesian literature (Tiao and Zellner (1964); Leamer (1972); Shiller (1973); Akaike (1979); Doan et al. (1984); De Mol et al. (2008); Goodwin et al. (1992); Banbura et al. (2010); Pillonetto and De Nicolao (2010); Pillonetto et al. (2011a); Chen et al. (2012)).

According to the discussion in the previous section, this assumption implicitly implies that the vector α is (a sub vector of) an exchangeable sequence. In our personal view this is indeed reasonable since the basis vectors ψ_i capture the properties of g while the coefficients α_i all have the same weight and play a symmetric role. Thus permutation invariance of their distribution seems a reasonable assumption. We shall leave it to the reader to judge whether this is an acceptable assumption to make or not, and indeed we would welcome comments and discussions on this.

Conversely, from (39), the exchangeability assumption on α only guarantees the existence of a random distribution Π , which needs not be Gaussian. Thus the hierarchical model (43) entails the additional assumption on the parametric form of Π (zero mean Gaussian with variance λ) so that (39) leads to (43).

Similarly to Aravkin et al. (2014) we now refer to model (41) and introduce the Singular Value Decomposition $Q D^{1/2} V^\top = \Phi \Psi$. Define also $Z := Q^\top Y$, $\beta := V^\top \alpha$ and with some abuse of notation denote $Q^\top E$ still as E (since $Q^\top E$ has the same distribution as E when $E \sim \mathcal{N}(0, \sigma^2 I)$).

The linear model (41) becomes

$$Z = D^{1/2} \beta + E = \bar{\beta} + E \quad \bar{\beta} := D^{1/2} \beta \quad (44)$$

A simple calculation shows that the Bayes estimator (say $\hat{\beta} = \delta(Z)$) (15) of g , in the "new" coordinates $\bar{\beta} = D^{1/2} V^\top \alpha$, takes the form of a shrinkage estimator; written componentwise:

$$\delta_i(Z) = \hat{\beta}_i = \left(1 - \frac{\sigma^2/d_i}{\lambda + \sigma^2/d_i}\right) Z_i \quad (45)$$

Note that model (44) has the same form as (29) in Aravkin et al. (2014) with the difference that now the prior covariance needs not be a multiple of the identity. It is now interesting to connect model (44) with the change of coordinates suggested in Strawderman (1978), which considers a linear estimation problem of the form (4), (41). We again here refer to Strawderman (1978), which introduces a change of coordinates in the parameter g (named θ in Strawderman (1978)), defining $\bar{g} = A g = A \Psi \alpha$ where A is such that⁵, for some diagonal matrix Λ ,

$$A K A^\top = \Lambda \quad A (\Phi^\top \Phi)^{-1} A^\top = I \quad (46)$$

⁵ Note that one should take B in Strawderman (1978) so that $B^{-1} := \Phi^\top \Phi K \Phi^\top \Phi$; with this definition of B and given A such that $A (\Phi^\top \Phi)^{-1} A^\top = I$ holds, $A K A^\top = \Lambda$ is equivalent to $A B A^\top = \Lambda^{-1}$.

The structure of this change of coordinates is characterized by the following lemma.

Lemma 4. Given the Singular Value Decompositions $\Phi\Psi = QD^{1/2}V^\top$ and $\Phi = PS^{1/2}U^\top$, the change of coordinates $\tilde{\beta} = Ag = A\Psi\alpha$ leading to (44) is such that (46) is satisfied with $\Lambda = D$ if

$$A = Q^\top PS^{1/2}U^\top = Q^\top \Phi = D^{1/2}V^\top \Psi^{-1} \quad (47)$$

Proof: It is a simple matrix manipulation to check that $A = \tilde{Q}S^{1/2}U^\top$ for some \tilde{Q} : $\tilde{Q}^\top \tilde{Q} = \tilde{Q}\tilde{Q}^\top = I$ satisfies $A(\Phi^\top \Phi)^{-1}A^\top = I$. Next, observe that

$$\begin{aligned} AK A^\top &= \tilde{Q}S^{1/2}U^\top KUS^{1/2}\tilde{Q}^\top \\ &= \tilde{Q}P^\top \left[PS^{1/2}U^\top KUS^{1/2}P^\top \right] P\tilde{Q}^\top \\ &= \tilde{Q}P^\top \left[\Phi K \Phi^\top \right] P\tilde{Q}^\top \\ &= \tilde{Q}P^\top QDQ^\top P\tilde{Q}^\top \end{aligned}$$

so that choosing $\tilde{Q} = Q^\top P$, $A = Q^\top PS^{1/2}U^\top$ satisfies (46) with $\Lambda = D$, the (squared) singular values of $\Phi\Psi$. \diamond

Theorem 1 in Strawderman (1978) provides, for model (44), a family of adaptive estimation rules which are *minimax* w.r.t. the loss

$$L(\tilde{\beta}, \delta) := (\delta - \tilde{\beta})^\top D (\delta - \tilde{\beta}) \quad (48)$$

obtained setting $Q = D$ in (38).

We shall now attempt to provide a link between the minimax rules given in Strawderman (1978) and the Empirical Bayes estimator obtained plugging $\hat{\lambda}_{ML}$ in (45).

Remark 5. Note that the expected value of (48) is exactly the weighted Mean Squared Error introduced and studied in Aravkin et al. (2014) adapted to the case in which the prior covariance is not limited to being a multiple of the identity. It was shown in Aravkin et al. (2014) that the marginal likelihood estimator $\hat{\lambda}_{ML}$ converges, as the number of data goes to infinity, to

$$\lambda^* := \frac{\|\beta\|^2}{B} = \frac{\tilde{\beta}^\top D^{-1} \tilde{\beta}}{B} = \frac{g^\top K^{-1} g}{B} \quad (49)$$

where the change of coordinates in Lemma 4 has been used. See also 4.2 (in particular eq. (18), (19) and Criterio 3 (page 863)) of Nicolao et al. (1997), vs (44), (45) and (49) of this paper.

It has been shown in Aravkin et al. (2014) that the estimator δ^* obtained setting $\lambda = \lambda^*$ in (45) also minimizes (asymptotically) the expected loss $\mathbb{E}L(\tilde{\beta}, \delta)$, thus

$$\mathbb{E}L(\tilde{\beta}, \delta^*) \leq \mathbb{E}L(\tilde{\beta}, \hat{\beta}_{LS})$$

where $\hat{\beta}_{LS} := Z$ is the least squares estimator of $\tilde{\beta}$. Now assume that, for N large, the marginal likelihood estimator is close to its limit, i.e.

$$\hat{\lambda}_{ML} \simeq \frac{\tilde{\beta}^\top D^{-1} \tilde{\beta}}{B} \simeq \frac{Z^\top D^{-1} Z}{B} = \frac{\|Z\|_{D^{-1}}^2}{B}$$

so that it makes sense to approximate the empirical Bayes estimator of $\tilde{\beta}$ replacing $\hat{\lambda}_{ML}$ with $\lambda = \frac{Z^\top D^{-1} Z}{B}$, leading to

$$\begin{aligned} \hat{\beta}_{EB,i} &= \left(1 - \frac{\sigma^2/d_i}{\hat{\lambda}_{ML} + \sigma^2/d_i} \right) Z_i \\ &\simeq \left(1 - \frac{B\sigma^2/d_i}{\|Z\|_{D^{-1}}^2 + B\sigma^2/d_i} \right) Z_i \end{aligned} \quad (50)$$

which has a form similar to equation (2.2) in Strawderman (1978). This discussion does not provide a formal proof that the estimator $\hat{\beta}_{EB}$ is minimax with respect to the loss (48), yet asymptotically it does dominate the least squares estimator as demonstrated in Aravkin et al. (2014).

The discussion above provides a clear connection between compound estimation problems studied in James and Stein (1961) and system identification. Indeed when performing System Identification, there is often no interest in the single basis expansion coefficients α_i in (40) or the impulse response coefficients g_i , but rather in their compound effect on output prediction error $\tilde{Y} := \Phi(g - \hat{g}) = \Phi A^{-1}(\tilde{\beta} - \delta)$.

Remark 6. The reader may argue that, sometimes, the interest is in the prediction w.r.t. a different input as that used for identification (e.g. low pass vs. white noise). This fact could be accounted for with an appropriate design of the kernel matrix K , see eq. (51) below.

Expressing the cost (48) in terms of output prediction error \tilde{Y} as:

$$\begin{aligned} L(\tilde{\beta}, \delta) &= L(\Phi g, \Phi \hat{g}) \\ &= \tilde{Y}^\top (\Phi K \Phi^\top) \tilde{Y} = \tilde{g}^\top (\Phi^\top \Phi K \Phi^\top \Phi) \tilde{g} \end{aligned} \quad (51)$$

it is seen that the loss (48) is a weighted output prediction error or, equivalently, a weighted ℓ_2 error on the impulse response fit. The weighting $\Phi K \Phi^\top = \Phi \Psi \Psi^\top \Phi^\top$ carries information about:

- (1) the “model class” specified through the choice of the basis vectors ψ_i (and thus the covariance matrix $K = \Psi \Psi^\top$)
- (2) how these vectors are mapped in output prediction through Φ , which of course depends on the excitation properties since the regressor matrix Φ contains the input $u(t)$.

This shows that the Bayes estimator equipped with the Marginal Likelihood estimator $\hat{\lambda}_{ML}$ attempts to minimize output prediction error along directions relevant to the model class considered (through the basis functions ψ_i) which are also properly “excited” by the input. As discussed in Aravkin et al. (2014) the weighted mean squared error (51) gives more emphasis to directions in the “output” space which have high signal to noise ratio. As a result the empirical Bayes estimator $\hat{\beta}_{EB}$, which (asymptotically) minimizes the weighted MSE (51), is robust w.r.t. noise in the measurements Y . We shall elaborate upon the consequences of this observation in the next subsection.

6.1 Robustness of the Marginal Likelihood Estimator

An important implication of this robustness has been discussed in Section 6 of Pillonetto and Chiuso (2015a) in connection with the concept of *excess degrees of freedom*. We recall here that the degrees of freedom of an estimator $\delta(Y)$ are defined (under Gaussian assumptions) in terms of the matrix of partial derivatives, see Remark (6.3) in Pillonetto and Chiuso (2015a),

$$[D_f(\delta)]_{ij} := \mathbb{E} \left[\frac{\partial (\delta(Y))_i}{\partial Y_j} \right] \quad (52)$$

$D_f(\delta)$ above is related to the so called *optimism* (Hastie et al. (2009)), i.e. the average amount by which the “fitting

error” underestimates the prediction error of the estimator on new data.

For instance, with reference to model (4) with $\hat{Y} := \Phi\hat{g} = \delta(Y)$, Y the “identification data” and Y_{test} some “fresh” test data, we have

$$\mathbb{E}\|Y_{test} - \Phi\hat{g}\|^2 = \mathbb{E}\|Y - \Phi\hat{g}\|^2 + 2\sigma^2 \text{Tr}\{D_f(\delta)\}$$

Since the estimator $\delta(Y) = \delta_{\hat{\lambda}(Y)}(Y)$ depends on the data also through the estimator $\hat{\lambda}(Y)$ of the “regularization” parameter $\lambda \in \mathbb{R}^k$, the degrees of freedom can be expressed as:

$$D_f(\delta) = \mathbb{E} \left[\frac{\partial \delta(Y)}{\partial Y} \right] + \sum_{i=1}^k \mathbb{E} \left[\frac{\partial \delta(Y)}{\partial \lambda_i} \frac{\partial \hat{\lambda}_i(Y)}{\partial Y} \right] \quad (53)$$

The trace of the rightmost term has been called *excess degrees of freedom* and accounts for the sensitivity of $\hat{\lambda}$ w.r.t. the data Y . Connection between optimism and sensitivities of a generic estimator is also studied in Dinuzzo et al. (2007), see Thm 2, page 2478.

Recall also that, for λ fixed, the a posteriori variance of $\hat{Y} := \Phi\hat{g} = \delta(Y)$ is related to the degrees of freedom by:

$$\text{Tr}[Var\{\Phi(g - \hat{g})\}] = \sigma^2 \text{Tr} \left(\mathbb{E} \left[\frac{\partial \delta(Y)}{\partial Y} \right] \right) = \text{Tr}[D_f(\delta)] \quad (54)$$

However the posterior variance of $\Phi\hat{g}$ when λ is estimated from data is larger than (54); in fact, also the uncertainty of $\hat{\lambda}$ should be accounted for as it (negatively) affects the posterior variance of the predictor. Therefore it is also desirable to control the variability of $\hat{\lambda}$ or, equivalently, the excess degrees of freedom. This is, to some extent, guaranteed by the robustness of the Marginal Likelihood estimator.

7. PRIORS FOR DYNAMICAL SYSTEMS: THE PRESENT

In Section 6 we started by postulating that the impulse response g is represented using the series expansion (40). This representation provides a connection with the extensive literature on ridge regression and minimax estimation in linear models. It has also been observed that (40) connects with Bayesian priors (43) provided $K := \sum_{i=1}^T \psi_i \psi_i^\top$. This latter decomposition can be, for instance, obtained from the Karhunen-Loève (SVD) decomposition of the kernel K . Thus, designing (or selecting) the basis functions ψ_i is equivalent to selecting the prior covariance matrix K . The recent literature in System Identification has extensively studied this problem from the Bayesian/Regularization point of view. The reader is also referred to Pillonetto et al. (2015) for connections with the so-called *atomic norm*.

In most of this paper we have drawn connections between the system identification problem (2) and estimation in the linear model (4). There are however important features of the unknown vector g in (4) which we may like to account for when performing identification. For instance, starting from Akaike (1979), (see also Section 4 for a thorough discussion) “priors” have been introduced which account for the fact that g should be a decreasing sequence (for instance in Akaike (1979) it is guaranteed that $\lim_{T \rightarrow \infty} \mathbb{E} \sum_{k=1}^T |g_k|^2 < \infty$). In this Section we shall

provide an overview of the recent developments along this direction. The reader is also referred to Dinuzzo (2012) for further discussions.

7.1 Priors for SISO systems

We shall now restrict to the simplest case $p = m = 1$. The MIMO case shall be discussed later on. In the seminal paper Pillonetto and De Nicolao (2010) the so called *Stable Spline* prior has been introduced. The basic building block is the family of spline priors (Wahba (1990), widely used in statistics and machine learning (Rasmussen and Williams (2006))); unfortunately these are not suitable to describe impulse responses of stable linear systems which are expected to be (exponentially) decaying sequences. The main novel idea, expanding upon a similar construction in Neve et al. (2007) (see e.g. page 1140), was to introduce an exponential time-warping

$$\tau = e^{-\beta t} \quad \beta > 0, \quad t \in \mathbb{R}^+ \quad (55)$$

on the family of spline priors so as to guarantee that realizations from the prior are, with probability one, exponentially decaying functions and thus represent BIBO stable linear systems. This reminds of the prior models discussed in Section 4 where different forms of decay rates for the variance of the coefficients g_i were postulated. It is fair to say that Goodwin et al. (1992) is probably the most closely related; in fact it was suggested that also the decay rate β (ρ in equation (24)) could be estimated from data using the marginal likelihood.

This is the way Stable Spline Kernels were introduced in Pillonetto and De Nicolao (2010); From an historical perspective, it is fair to say that the Stable Spline kernels are indeed mixing ideas from Shiller (1973), where *time-domain* smoothness conditions were advocated for the impulse response, and Akaike (1979) where a *frequency-domain* smoothness penalty was introduced. In fact, as also discussed in Goodwin et al. (1992), the exponential decay of the impulse response variances, implies a Lipschitz condition on the frequency response function (26).

The simplest version of these kernels is the so-called “first-order” stable spline prior, known also as Tuned-Correlated (TC) kernel, see Chen et al. (2012). It can be seen as the covariance sequence of the sampled, time-warped, Wiener process, so that the matrix K has entries

$$[K_{\lambda, \beta}]_{i,j} = \mathbb{E} g_i g_j = \lambda \min(e^{-\beta i T_c}, e^{-\beta j T_c}) \quad (56)$$

where T_c is the sampling time of the discrete time system.

Using the Karhunen-Loève decomposition (random Fourier series) of the Wiener process (see e.g. Anderson and Stephens (1997), eq. (11)) it is a simple calculation to see that g admits the series expansion

$$g_k = \alpha_0 e^{-\beta k T_c} + \sqrt{2} \sum_{i=1}^{\infty} \alpha_i \frac{\sin(\pi i e^{-\beta k T_c})}{\pi i} \quad (57)$$

$$\alpha_k \sim \mathcal{N}(0, \lambda) \quad i.i.d.$$

which directly connects with (40) provided one defines

$$\psi_0 = [1 \ e^{-\beta T_c} \ e^{-\beta 2 T_c} \ \dots \ e^{-\beta (T-1) T_c}]^\top$$

$$\psi_i = \sqrt{2} \left[\frac{\sin(\pi i)}{\pi i} \ \frac{\sin(\pi i e^{-\beta T_c})}{\pi i} \ \dots \ \frac{\sin(\pi i e^{-\beta (T-1) T_c})}{\pi i} \right]^\top \quad (58)$$

Note that the vectors ψ_i are not orthogonal (the time-warping (55) destroys the orthogonality of the Karhunen-Loeve basis) and indeed working directly with the SVD of the matrix (56) could be advantageous also from a computational point of view (Carli et al. (2012); Chen and Ljung (2013)). Yet the expansion (57) with basis (58) shows that the impulse response g is modeled by the Stable Spline prior as a decaying exponential $e^{-\beta k T_c}$ (ψ_0) plus a random perturbation which is a time-warped Brownian Bridge (see Chiuso et al. (2014)). As discussed in Chiuso et al. (2014), also the time-warped Brownian Bridge has a local behavior (as $k \rightarrow +\infty$) of the form $e^{-\beta k T_c}$, which derives from the fact that $\sin(x) \simeq x$ for $x \rightarrow 0$. This suggests that TC kernels, and possibly their conic combinations (Chen et al. (2014), Pillonetto et al. (2015)), are well suited to describe impulse responses of finite dimensional linear systems.

There is an interesting maximum entropy interpretation of the stable spline kernels, which can be derived under suitable constraints on the exponential decay of derivatives of g (in the continuous time formulation), see De Nicolao et al. (1998); Chen et al. (2015) and references therein. This very same property can be easily obtained observing that the Stable-Spline Kernels can be thought of as the covariances of time-varying backward AR processes. For instance for the TC/Ctable spline of order 1 kernel (56), it is an easy exercise to show that it can be obtained, on any finite interval $k \in [0, T]$, as the covariance of the AR(1) process (assuming w.l.o.g. $T_c = 1$)

$$\begin{aligned} g_T &\simeq \mathcal{N}(0, c\lambda^T) & g_T \perp n_i \quad \forall i \\ g_{k-1} &= g_k + \sqrt{c\lambda^{k-1}(1-\lambda)}n_k & n_k \simeq \mathcal{N}(0, 1) \quad i.i.d. \end{aligned} \quad (59)$$

This simple observation allows to derive its Maximum Entropy properties as well as the band structure of the inverse covariance directly from well known results on AR processes (Grenander and Szegő (1958)).

Several variations of (56) have been proposed, including higher order splines (Pillonetto and De Nicolao (2010)) and their filtered versions (Pillonetto et al. (2011a)). Chen et al. (2012) also introduces other possibilities (notably the so called DC kernels) while later papers (Chen et al. (2014); Chiuso et al. (2014)) investigate the potential advantages, both from the point of computational efficiency as well as in terms of their flexibility, of building multiple kernels combining “elementary” kernels (TC, Stable Splines, DC etc.).

A common feature of all these kernels is that they all depend upon at least two hyperparameters (β and λ) and very often more are needed. Let us say that the hyperparameter vector is partitioned in λ, ξ . The first is a “scale” parameter (the variance of the coefficients α_i) which, as we shall see later, can also be used to perform variable selection. The second one, ξ , can be seen as a “shape” parameter vector which controls the shape of the basis functions ψ_i . For the TC kernel discussed above, $\xi = \beta$. The analysis reported in Section 6 relies on the fact that ξ is fixed, while in practice it needs to be tuned (e.g. via marginal likelihood maximization (17) or, e.g., cross validation). Much work is needed to fully clarify this

issue, see Pillonetto and Chiuso (2015a) for some results in support of the use of the marginal likelihood.

A complementary approach has been advocated in Darwish et al. (2014); Chen and Ljung (2015) where the authors started directly from an orthonormal basis (e.g. the Laguerre basis, Wahlberg (1991)) and then placed a (decaying) prior on the coefficients (see eqs. (23) and (24) in Chen and Ljung (2015)) to guarantee convergence of the series (see (42)); This leads to an overall Kernel $K_{\lambda, \xi}$ which depends on the parameters ξ describing the basis functions (called p in Chen and Ljung (2015)) and those describing the prior on the coefficients (called α in Chen and Ljung (2015)); it is suggested that (λ, ξ) are estimated using marginal likelihood optimization. The results in Chen and Ljung (2015) suggest that this kernel may have some advantages w.r.t. the TC/Stable Spline kernels discussed above.

More recent developments (see e.g. Prando et al. (2014)) aim at building Kernels which favour estimators \hat{g} with low McMillan degree. This is done imposing a suitable penalty on the Hankel matrix of the estimated impulse response. A formulation in the framework of Gaussian priors will be discussed in Section 8.

7.2 Priors for MIMO systems

When the input and/or output data have dimension larger than one (i.e. $p > 1$ and/or $m > 1$) the impulse response coefficients g_k are $p \times m$ real matrices. The simplest solution, advocated in Chiuso and Pillonetto (2012), is to consider the impulse responses $([g_1]_{ij}, \dots, [g_T]_{ij})$ from input j to output i as independent random vectors. Therefore, if the matrix sequence g_k has been vectorized component wise, i.e.

$$g = \left([g_1]_{11} \dots [g_T]_{11}, [g_1]_{21} \dots [g_T]_{21} \dots [g_1]_{pm} \dots [g_T]_{pm} \right)^\top$$

the prior covariance K has a block diagonal structure

$$K_{\lambda, \beta} = \text{block diag} \{ K_{\lambda_{11}, \beta_{11}}, \dots, K_{\lambda_{pm}, \beta_{pm}} \} \quad (60)$$

where

$$\lambda := [\lambda_{11}, \dots, \lambda_{pm}] \quad \beta := [\beta_{11}, \dots, \beta_{pm}]$$

Depending on the problem one may assume that $\beta_j = \beta_i$, $\forall i, j$ so that all kernels share the same decay rate. Instead the scale factors λ_i are typically all different and this also serves to the purpose of performing variable selection, see Section 8 for more details.

Additional issues concerning “structure”, such as low McMillan degree, presence of latent variables and so on, will be discussed in the next section.

8. STRUCTURE IN MODERN SYSTEM IDENTIFICATION AND SPARSITY

In addition to stability of the impulse response g , in System Identification one is often interested in favoring (and discovering) a certain type of hidden structure in the identified model, the more so as the number of variables (and thus model “size”) grows w.r.t. the number of available data. This structure may have to do with model complexity (measured by McMillan degree), the presence of latent (unmeasurable) factors (see e.g. Forni et al. (2004)), or in terms of Granger causality structure

(see Granger (1963)) which as we shall see in Section 8.1 is related to sparsity; we refer to these general properties as *structure*.

In the traditional “parametric” framework system identification has been performed as the cascade of model (and hence structure) selection followed by parameter estimation. Regularization approaches permit to perform this structure selection directly as part of the “estimation” process. The unknown g is embedded in a “high dimensional” space and, eventually, the estimator \hat{g} is pushed towards a low dimensional manifold which is to be learned from data. This manifold may be that of low McMillan degree models, sparse impulse response matrices, or models admitting latent variables representations and so on. As it will become clear finding this underlying structure from data can be formulated as a sparsity problem.

8.1 Variable Selection in System Identification

Consider a MIMO OE system (2) with transfer matrix⁶ $G(q)$; the matrix $G(q)$ may have so zero entries (meaning that $G_{ij}(q) = 0$), so that the i – th output does not depend on the j – th input. Detecting the zero entries of $G(q)$ from data corresponds to performing variable selection, or equivalently finding the structure of a causality graph which describes the “interaction” between inputs and outputs; this problem has been framed in the context of sparsity in Materassi et al. (2009); Chiuso and Pillonetto (2010b); Songsiri and Vandeberghe (2010); Sanandaji et al. (2011), where the link with literature on sparse linear models (Tibshirani (1994); Efron et al. (2004)) and compressive sensing (Donoho (2006)) has been exploited. Among the papers mentioned above, only Chiuso and Pillonetto (2010b) explicitly introduces regularization on the impulse responses to avoid overfitting when the number of lags T needs to be large. Further developments can be found in Chiuso and Pillonetto (2010a, 2012) where Bayesian techniques have been proposed. In fact this “Bayesian” approach (which corresponds to using the kernel (60) in conjunction with marginal likelihood estimators for β and λ) is related to so-called Automatic Relevance Determination (ARD, Mackay (1994)) and Sparse Bayesian Learning (SBL, Tipping (2001)), well known in the Machine Learning community. In fact, as shown in Chiuso and Pillonetto (2012); Chen et al. (2014), the marginal likelihood estimators $\hat{\lambda}_{ij}$ of λ_{ij} in (60) are “threshold estimators” meaning that $\hat{\lambda}_{ij} = 0$ on a set of data of positive measure.

Example 1. A very simple example is the following: assume data y_i , $i = 1, \dots, N$ are generated according to

$$y_i = \theta + e_i \quad \theta \in \mathbb{R}$$

where e_i are i.i.d. zero mean Gaussian variables with known variance σ^2 . Assuming a prior $\theta \sim \mathcal{N}(0, \lambda)$, it is an easy exercise to show that the maximum marginal likelihood estimator $\hat{\lambda}_{ML}$ (17) is given by (see also eq. (35)):

$$\hat{\lambda}_{ML} = \max \left(0, \left(\frac{1}{N} \sum_{i=1}^N y_i \right)^2 - \sigma^2 \right)$$

showing that $\hat{\lambda}_{ML}$ is non negative only provided the squared sample average $\left(\frac{1}{N} \sum_{i=1}^N y_i \right)^2$ of y is above the noise variance σ^2 .

When $\hat{\lambda}_{ij} = 0$ the prior for $([g_1]_{ij} \dots [g_T]_{ij})$, the impulse response from input j to output i , has zero mean and zero variance resulting in a posterior estimator $\mathbb{E}_{\hat{\lambda}}[[g_k]_{ij} | Y] = 0$, $\forall k = 1, \dots, T$. There is evidence (see e.g. Chiuso and Pillonetto (2012) in the context of dynamical systems) that these Bayesian approaches are superior to “Lasso-type” procedures, and in fact an extensive literature has attempted with some success to show this, see Wipf et al. (2011); Aravkin et al. (2014) and references therein. In particular Aravkin et al. (2014) shows that these Bayesian procedures have a better “shrinking vs. sparsity” tradeoff than “Lasso-type” procedures (such as Group Lasso).

8.2 Low McMillan degree, Hankel matrices and sparsity

A major limitation of a Gaussian prior with block diagonal covariance K as in (60) is that this prior does not capture the inherent relations between the transfer functions $G_{ij}(q)$ when $G(q)$ has low McMillan degree. This is particularly important if m and p are large. Several authors have addressed this issue. For instance Fazel et al. (2001) tackles this problem replacing $J_\eta(g)$ in (11) with a nuclear norm penalty,

$$J_\eta(g) := \eta \|\mathcal{H}(g)\|_* = \eta \text{Trace} \sqrt{\mathcal{H}(g)\mathcal{H}(g)^\top} \quad (61)$$

where $\mathcal{H}(g)$ is the block Hankel matrix built with the impulse response sequence g . The reason for using the nuclear norm stems from Theorem 1 in Fazel et al. (2001) which shows that the nuclear norm $\|A\|_*$ is the convex envelope of the rank of A on the ball $\|A\| < 1$. Since $J_\eta(g)$ approximates the rank function it is expected that solutions \hat{g} of (11) with low nuclear norm $\|\mathcal{H}(\hat{g})\|_*$ are close to having $\mathcal{H}(\hat{g})$ of low rank and thus low McMillan degree of \hat{g} . Tens of papers have been published on the use of the nuclear norm in System Identification, and it is certainly impossible to survey them here. We shall only mention Shah et al. (2012) which discusses the link between Hankel nuclear norm and *atomic* norm, Hjalmarrsson et al. (2012) which applies these ideas to Box-Jenkins modeling, Chiuso et al. (2013) where Nuclear Norm-based techniques are first compared in a simulation study to other regularization methods, Liu and Vandenberghe (2009), Ayazoglu and Sznajder (2012) and Blomberg et al. (2015) where computational issues were addressed, and Verhaegen and Hansson (2015) which applies nuclear norm penalties to subspace-type algorithms. See also Prando et al. (2015) for an overview.

Unfortunately, using the nuclear norm penalty is known to introduce a (possibly significant) bias in the estimator (11). This can be easily explained by the fact that the nuclear norm is really an ℓ_1 -norm on the singular values of the Hankel matrix, and thus related to Lasso-type approaches. Lasso is well known to introduce significant bias on large coefficients. This is the reason why variations of the Lasso (such as the Adaptive Lasso Zou (2006) or SCAD Fan and Li (2001) and so on) have been introduced in the Statistics literature.

⁶ The same type of considerations hold for a general multivariable ARMAX model (1), see Chiuso and Pillonetto (2012).

Inspired by the analysis of the “Shrinking vs. Sparsity” tradeoff of Lasso and Bayesian procedures found in Aravkin et al. (2014) and by some recent papers on rank minimization (see Wipf (2012); Mohan and Fazel (2012)), the paper Prando et al. (2014) introduces a new “Bayesian” procedure for rank minimization in the context of system identification.

A “general” form (see Wipf (2012); Prando et al. (2014)) of a low-rank promoting prior for a matrix A , which involves quadratic forms of A , has the structure

$$p_\lambda(A) \propto e^{-\lambda \text{Tr}(QAA^\top)} \quad (62)$$

where $Q = Q^\top \geq 0$ plays the role of an hyperparameter. Slightly different structures have been proposed in Sundin et al. (2014) where “left” and “right” hyperparameters have been considered.

Here, as in Prando et al. (2014), we consider a simplified version of (62) which can be derived under maximum entropy arguments (see Prando et al. (2015)), where the matrix Q takes on a very simple structure

$$Q = \lambda_1 \Pi_{U_n} + \lambda_2 \Pi_{U_n^\perp};$$

Π_{U_n} and $\Pi_{U_n^\perp}$ are, respectively, the projectors on the column spaces of U_n and on its orthogonal complement U_n^\perp . This, complemented with a “stable spline” prior yields to the so called *Stable-Hankel* (SH) prior

$$p_\eta(g) \propto e^{-\frac{1}{2}J_\eta(g)} \\ J_\eta(g) := \lambda_s g^\top K_s^{-1} g + \lambda_1 \text{Trace}[\Pi_{U_n} \mathcal{H}(g) \mathcal{H}(g)^\top] + \lambda_2 \text{Trace}[\Pi_{U_n^\perp} \mathcal{H}(g) \mathcal{H}(g)^\top] \quad (63)$$

where K_s is a “stable” kernel (with decaying diagonal elements) so as to guarantee the estimators \hat{g} are decaying functions of the lag. The recent paper Pillonetto et al. (2015) discusses the role of the stability constraint, and the effect of not exploiting it in Hankel nuclear norm based penalties.

The scale factors λ_s , λ_1 and λ_2 are estimated using the marginal likelihood (17); also in this case the marginal likelihood can be computed in closed form, being (63) a quadratic form in g . Also the matrix U_n along with the number of its columns n can be estimated using the marginal likelihood, see Prando et al. (2015) for details. Suffices here to observe that the splitting of the penalty along the (column) spaces of U_n and its orthogonal complement U_n^\perp allows to reduce the bias along the “signal” subspace, while being robust w.r.t. “errors” within $\text{col span}\{U_n\}$. This research direction is still in its infancy and much work is needed to fully clarify the potential and limitations of this approach. The interested reader is referred to Prando et al. (2015) which has been recently presented at the 2015 IFAC symposium on System Identification.

8.3 Latent variable models and sparsity

When dealing with high dimensional data, and thus predictors depending upon a large number of variables, there is a need for parsimonious models which are at the same time robust (w.r.t. noise in the data) and interpretable. As an example consider the linear static model

$$y = Mx + e$$

where the data y are expressed as a linear function Mx of some variable x (measured or unmeasured). Two common alternatives explored in the literature to control the

complexity of M is to either assume it is of *low rank* ($M = M_{LR}$) or it is *sparse* ($M = M_S$), or combinations thereof ($M = M_{LR} + M_S$, see e.g. Chandrasekaran et al. (2011) and references therein for the use of this decomposition in a static scenario).

The low rank component M_{LR} models the presence of few underlying “common” variables which capture the *co-movement* of the observables y . Relevant keywords here are Principal Component Regression and Factor Models. On the other hand a sparse model can be encoded in terms of a network of (conditional) dependences (Lauritzen (1996)).

Such models can also be studied in a dynamic scenario where y is a multivariate stochastic process and the matrix M is replaced by a transfer matrix $G(q)$. For instance the alternative between considering Principal Component Regression and Regularized Regression when the number of “regressors” becomes large is studied in De Mol et al. (2008).

However a through study of “Sparse + Low rank” dynamical models seems to be lacking; the paper Zorzi and Sepulchre (2014) discusses Sparse plus Low rank AR models where sparsity is related to smoothing conditions in the style of Songsiri and Vandeberghe (2010). If instead causal (*a la Granger*) conditional dependences are considered, some preliminary work in a Bayesian framework can be found in (Zorzi and Chiuso (2015)) where the output process $y(t) \in \mathbb{R}^p$ is modeled as

$$\begin{aligned} y(t) &= Fx(t) + S(q)y(t) + v(t) \\ x(t) &= H(q)y(t) + w(t) \end{aligned} \quad (64)$$

where $x(t) \in \mathbb{R}^r$, $r \ll p$ is a latent process (not measurable), $S(q)$ and $H(q)$ are strictly causal stable transfer matrices and $S(q)$ is sparse in the sense that many of its entries $S_{ij}(q)$ are null transfer functions. Model (64) can be given a graphical interpretation with the components of $y(t)$ being “visible” nodes, those of $x(t)$ being hidden nodes and edges encoding the (Granger causality) dependences between these time series. The Bayesian framework discussed in this paper can be used to simultaneously infer the transfer matrices $S(q)$, $H(q)$, the matrix F as well as the number of latent variables $x(t)$. Details can be found in Zorzi and Chiuso (2015).

9. CONCLUSIONS AND FUTURE OUTLOOK

As briefly outlined in Sections 7 and 8 the recent literature has addressed the problem of designing Kernel structure which can capture (and actually favour) structural properties of dynamical systems. In a sense the Bayesian formulation can be thought of as *lifting one level up* the classical problem of estimating the system parameters, to the problem of determining the parameters of a flexible model class (the hyperparameters describing the prior), see also Lindley and Smith (1972).

Of course there is a tradeoff between how well a prior density can be shaped to give low bias and how well the hyperparameters can be estimated from data (low variance), thus leading to the usual *bias-variance tradeoff*, which can be faced again via a *parsimony* principle. We believe much work remains to be done along this direction,

both in terms of design as well as in terms of analysis of the tuning procedures.

In terms of kernel design, it should be apparent that the problems considered in this paper are only a few examples (possibly the most straightforward ones) and it is indeed to be expected that many others can be considered and framed with the Bayesian language. Yet the reader should keep in mind that the System Identification problem should not be regarded as a general function regression (or interpolation) problem. There are intrinsic properties of dynamical systems that should always be accounted for and used when designing prior models or, equivalently, model classes.

As the dimension of data increases, both in terms of number of time series p which need to be jointly modeled, as well as in terms of number of observations N one is able or should process at a time, the computational aspect becomes more and more relevant. The more so if real time use of these modeling tools is envisaged. For instance in Bonettini et al. (2015) fast first order methods are considered for marginal likelihood optimization. These techniques have been adapted and used in Prando et al. (2015) showing significant improvements in terms of computational time w.r.t. off-the-shelf techniques.

Yet much work needs to be done to make the techniques discussed in this paper scale well with p and N ; to such purpose we believe that optimization issues should enter at a much earlier stage (such as the design of the prior) so as to guide the developments of algorithms and methods. We envision that stochastic optimization and possibly early stopping ideas may turn out to be useful in this context (see e.g. Rosasco et al. (2014)).

We also believe that further work is needed regarding the use, design and implementation of MCMC methods for handling either more complex situations, e.g. where the marginal likelihood cannot be computed in closed form Pillonetto et al. (2011b), or for computing so-called “full-Bayes” solutions and accurately describing the uncertainty in the estimators (Neve et al. (2008), Romeres et al. (2015)), not necessarily properly quantified by the empirical Bayes estimators, Pillonetto and Chiuso (2015a).

ACKNOWLEDGEMENTS

The author is grateful to a number of colleagues with whom he has shared ideas, thoughts and joint research efforts in the past years as well as for providing comments on early versions of this manuscript. Above all my colleague Gianluigi Pillonetto, coauthor of a long list of papers, but also, in alphabetical order, Tianshi Chen, Manfred Deistler, Giuseppe De Nicolao, Michel Gevers, Lennart Ljung, Giorgio Picci, Victor Solo, Roberto Tempo, Mattia Zorzi and my present students Giulia Prando and Diego Romeres.

REFERENCES

Akaike, H. (1974). A new look at the statistical model identification. *IEEE Trans. on Automatic Control*, AC-19, 716–723.

- Akaike, H. (1979). Smoothness priors and the distributed lag estimator. Technical report, Department of Statistics, Stanford University.
- Anderson, T. and Stephens, M. (1997). The continuous and discrete Brownian bridges: Representations and applications. *Linear Algebra and its Applications*, 264(0), 145 – 171.
- Aravkin, A., Burke, J., Chiuso, A., and Pillonetto, G. (2014). Convex vs non-convex estimators for regression and sparse estimation: the mean squared error properties of ARD and GLasso. *Journal of Machine Learning Research*, 15, 1–36.
- Aronszajn, N. (1950). Theory of reproducing kernels. *Trans. of the American Mathematical Society*, 68, 337–404.
- Åström, K. and Bohlin, T. (1965). Numerical Identification of Linear Dynamic Systems from Normal Operating Records. In *IFAC Symposium on Self-adaptive Control Systems*.
- Ayazoglu, M. and Sznajder, M. (2012). An algorithm for fast constrained nuclear norm minimization and applications to systems identification. In *IEEE CDC 2012, December 10-13, 2012, Maui, HI, USA*, 3469–3475.
- Banbura, M., Giannone, D., and Reichlin, L. (2010). Large Bayesian VARs. *Journal of Applied Econometrics*, 25(1), 71–92.
- Berger, J. (1982). Selecting a minimax estimator of a multivariate normal mean. *The Annals of Statistics*, 10, 81–92.
- Bertero, M. (1989). Linear inverse and ill-posed problems. *Advances in Electronics and Electron Physics*, 75, 1–120.
- Blomberg, N., Rojas, C., and Wahlberg, B. (2015). Approximate regularization paths for nuclear norm minimization using singular value bounds – with implementation and extended appendix. Technical report, arXiv:1504.05208.
- Bonettini, S., Chiuso, A., and Prato, M. (2015). A scaled gradient projection methods for Bayesian learning in dynamical systems. *SIAM Journal on Scientific Computing*, in press.
- Box, G., Jenkins, G., and Reinsel, G. (2008). *Time Series Analysis: Forecasting & Control*. Wiley, 4rd edition.
- Brockwell, P. and Davis, R. (1991). *Time Series, Theory and Methods*. Springer, Berlin.
- Carli, F., Chiuso, A., and Pillonetto, G. (2012). Efficient algorithms for large scale linear system identification using stable spline estimators. In *Proc. of SYSID*.
- Casella, G. (1980). Minimax ridge regression estimation. *The Annals of Statistics*, 8, 1036–1056.
- Chandrasekaran, V., Sanghavi, S., Parrilo, P.A., and Willsky, A.S. (2011). Rank-sparsity incoherence for matrix decomposition. *SIAM Journal on Optimization*, 21(2), 572–596.
- Chen, T., Andersen, M., Ljung, L., Chiuso, A., and Pilonetto, G. (2014). System identification via sparse multiple kernel-based regularization using sequential convex optimization techniques. *IEEE Trans. on Automatic Control*, 56(11), 2933–2945.
- Chen, T., Ardeschiri, T., Carli, F., Chiuso, A., Ljung, L., and Pilonetto, G. (2015). Maximum entropy properties of discrete-time first-order stable spline kernel. Technical report, arXiv:1504.03253.

- Chen, T. and Ljung, L. (2013). Implementation of algorithms for tuning parameters in regularized least squares problems in system identification. *Automatica*, 49(7), 2213–2220.
- Chen, T. and Ljung, L. (2015). Regularized system identification using orthonormal basis functions. Technical report, arXiv:1504.02872.
- Chen, T., Ohlsson, H., and Ljung, L. (2012). On the estimation of transfer functions, regularizations and Gaussian processes - revisited. *Automatica*, 48(8), 1525–1535.
- Chiuso, A., Chen, T., Ljung, L., and Pillonetto, G. (2013). Regularization strategies for nonparametric system identification. In *Proc. of IEEE CDC*.
- Chiuso, A., Chen, T., Ljung, L., and Pillonetto, G. (2014). On the design of multiple kernels for nonparametric linear system identification. In *IEEE CDC 2014*.
- Chiuso, A. and Pillonetto, G. (2010a). Learning sparse dynamic linear systems using stable spline kernels and exponential hyperpriors. In *Proceedings of Neural Information Processing Symposium*. Vancouver.
- Chiuso, A. and Pillonetto, G. (2010b). Nonparametric sparse estimators for identification of large scale linear systems. In *IEEE CDC 2010*, 2942–2947.
- Chiuso, A. and Pillonetto, G. (2012). A Bayesian approach to sparse dynamic network identification. *Automatica*, 48(8), 1553 – 1565.
- Cramér, H. (1946). *Mathematical Methods of Statistics*. Almqvist & Wiksells Akademiska Handböcker. Princeton University Press.
- Darwish, M., Toth, R., and Van den Hof, P. (2014). Bayesian system identification based on generalized orthonormal basis functions. Technical report, ERNSI Workshop.
- de Finetti, B. (1931). La funzione caratteristica di un fenomeno aleatorio.
- De Mol, C., Giannone, D., and Reichlin, L. (2008). Forecasting using a large number of predictors: Is Bayesian shrinkage a valid alternative to principal components? *Journal of Econometrics*, 146(2), 318 – 328.
- De Nicolao, G., Ferrari Trecate, G., and Lecchini, A. (1998). MaxEnt priors for stochastic filtering problems. In *Proc. Symp. on Math. Theory of Networks and Systems (MTNS’98)*, 755–758. Padova, Italy.
- Dinuzzo, F., Neve, M., Nicolao, G.D., and Gianazza, U. (2007). On the representer theorem and equivalent degrees of freedom of SVR. *Journal of Machine Learning Research*, 8, 2467–2495.
- Dinuzzo, F. (2012). Kernels for linear time invariant system identification. Technical report, arXiv:1203.4930.
- Doan, T., Litterman, R., and Sims, C. (1984). Forecasting and conditional projection using realistic prior distributions. *Econometric Reviews*, 3, 1–100.
- Donoho, D. (2006). Compressed sensing. *IEEE Trans. on Information Theory*, 52(4), 1289–1306.
- Edelman, D. (1988). Estimation of the mixing distribution for a normal mean with applications to the compound decision problem. *Ann. Statist.*, 16(4), 1609–1622.
- Efron, B., Hastie, T., Johnstone, L., and Tibshirani, R. (2004). Least angle regression. *Annals of Statistics*, 32, 407–499.
- Efron, B. and Morris, C. (1973). Stein’s estimation rule and its competitors—an empirical Bayes approach. *Journal of the American Statistical Association*, 68(341), 117–130.
- Fan, J. and Li, R. (2001). Variable selection via nonconcave penalized likelihood and its oracle properties. *Journal of the American Statistical Association*, 96(456), 1348–1360.
- Fazel, M., Hindi, H., and Boyd, S.P. (2001). A rank minimization heuristic with application to minimum order system approximation. In *Proc. of the 2001 ACC*, 4734–4739.
- Forni, M., Hallin, M., Lippi, M., and Reichlin, L. (2004). The generalized dynamic-factor model: identification and estimation. *The Review of Economics and Statistics*, 82, 540–554.
- Foster, M. (1961). An application of the Wiener-Kolmogorov smoothing theory to matrix inversion. *Journal of the Society for Industrial and Applied Mathematics*, 9(3), pp. 387–392.
- Giannone, D., Lenza, M., and Primiceri, G.E. (2015). Prior selection for vector autoregressions. *Review of Economics and Statistics*, 97(2), 436–451.
- Good, I. (1965). *The Estimation of Probabilities*. Research Monographs No. 30. MIT Press.
- Goodwin, G., Gevers, M., and Ninness, B. (1992). Quantifying the error in estimated transfer functions with application to model order selection. *IEEE Transactions on Automatic Control*, 37(7), 913–928.
- Goodwin, G.C. and Salgado, M.E. (1989). A stochastic embedding approach for quantifying uncertainty in the estimation of restricted complexity models. *International Journal of Adaptive Control and Signal Processing*, 3(4), 333–356.
- Granger, C. (1963). Economic processes involving feedback. *Information and Control*, 6, 28–48.
- Grenander, U. (1950). *Stochastic Processes and Statistical Inference*. Arkiv för matematik. Almqvist & Wiksell.
- Grenander, U. and Szegő, G. (1958). *Toeplitz Forms and their Applications*. Chelsea.
- Hannan, E. and Deistler, M. (1988). *The Statistical Theory of Linear Systems*. Wiley.
- Hastie, T.J., Tibshirani, R.J., and Friedman, J. (2009). *The Elements of Statistical Learning. Data Mining, Inference and Prediction*. Springer, Canada.
- Hewitt, E. and Savage, L.J. (1955). Symmetric measures on Cartesian products. *Trans. Amer. Math. Soc.*, 80(1), 470–501.
- Hjalmarsson, H., Welsh, J., and Rojas, C. (2012). Identification of box-jenkins models using structured ARX models and nuclear norm relaxation. In *Proc. of SYSID*, 322–327.
- Hoerl, A.E. and Kennard, R.W. (1970). Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics*, 12(1), 55–67.
- Hoerl, A. (1962). Application of ridge analysis to regression problems. *Chemical Engineering Progress*, 54–59.
- Hurvich, C. and Tsai, C. (1989). Regression and time series model selection in small samples. *Biometrika*, 76, 297–307.
- James, W. and Stein, C. (1961). Estimation with quadratic loss. In *Proc. 4th Berkeley Sympos. Math. Statist. and Prob., Vol. I*, 361–379. Univ. California Press, Berkeley, Calif.

- Johansen, T. (1997). On Tikhonov regularization, bias and variance in nonlinear system identification. *Automatica*, 33, 441–446.
- Kitagawa, G. and Gersch, W. (1996). *Smoothness Priors Analysis of Time Series*. Ima Volumes in Mathematics and Its Applications. Springer New York.
- Kitagawa, G. and Gersh, H. (1984). A smoothness priors-state space modeling of time series with trends and seasonalities. *Journal of the American Statistical Association*, 79(386), 378–389.
- Kitagawa, G. and Gersh, H. (1985). A smoothness priors long AR model methods for spectral estimation. *IEEE Transactions on Automatic Control*, 30(1), 57–65.
- Knapik, B.T., van der Vaart, A.W., and van Zanten, J.H. (2011). Bayesian inverse problems with Gaussian priors. *Ann. Statist.*, 39(5), 2626–2657.
- Knox, T., Stock, J., and Watson, M. (2001). Empirical Bayes forecast of one time series using many predictors. Technical report, National Bureau of Economic Research.
- Lauritzen, S. (1996). *Graphical Models*. Oxford University Press, Oxford.
- Leamer, E.E. (1972). A class of informative priors and distributed lag analysis. *Econometrica*, 40(6), pp. 1059–1081.
- Leeb, H. and Pötscher, B.M. (2005). Model selection and inference: Facts and fiction. *Econometric Theory*, 21(01), 21–59.
- Leeb, H. and Pötscher, B.M. (2006). Can one estimate the conditional distribution of post-model-selection estimators? *Ann. Statist.*, 34(5), 2554–2591.
- Lehmann, E.L. and Romano, J.P. (2005). *Testing statistical hypotheses*. Springer Texts in Statistics. Springer, New York, third edition.
- Lindley, D. and Smith, A. (1972). Bayes estimates for the linear model. *Journal of the Royal Statistical Society, Series B (Methodological)*, 34(1), 1–41.
- Liu, Z. and Vandenberghe, L. (2009). Interior-Point Method for Nuclear Norm Approximation with Application to System Identification. *SIAM Journal on Matrix Analysis and Applications*, 31(3), 1235–1256.
- Ljung, L. (1999). *System Identification - Theory for the User*. Prentice-Hall, Upper Saddle River, N.J., 2nd edition.
- Ljung, L., Goodwin, G., and Aguero, J. (2014). Stochastic embedding revisited: A modern interpretation. In *IEEE CDC*, 3340–3345.
- Ljung, L., Hjalmarsson, H., and Ohlsson, H. (2011). Four encounters with system identification. *European Journal of Control*, 17(5-6), 449 – 471.
- Loève, M. (1963). *Probability Theory*. Van Nostrand Reinhold.
- Lütkepohl, H. (2007). *New Introduction to Multiple Time Series Analysis*. Springer Publishing Company, Incorporated.
- Mackay, D. (1994). Bayesian non-linear modelling for the prediction competition. *ASHRAE Trans.*, 100(2), 3704–3716.
- Maritz, J.S. and Lwin, T. (1989). *Empirical Bayes Method*. Chapman and Hall.
- Marquardt, D.W. and Snee, R.D. (1975). Ridge regression in practice. *The American Statistician*, 29(1), pp. 3–20.
- Maruyama, Y. and Strawderman, W. (2005). A new class of generalized Bayes minimax ridge regression estimators. *The Annals of Statistics*, 1753–1770.
- Materassi, D., Innocenti, G., and Giarre, L. (2009). Reduced complexity models in the identification of dynamical networks: Links with sparsification problems. In *Proceedings of the 48th IEEE CDC 2009 (joint with IEEE CCC2009)*, 4796–4801.
- Melkman, A. and Micchelli, C. (1979). Optimal estimation of linear operators in hilbert spaces from inaccurate data. *SIAM Journal on Numerical Analysis*, 16, 87–105.
- Mohan, K. and Fazel, M. (2012). Iterative reweighted least-squares for matrix rank minimization. *Journal of Machine Learning Research*, 13, 3441–3473.
- Neve, M., De Nicolao, G., and Marchesi, L. (2008). Nonparametric identification of population models: An mcmc approach. *Biomedical Engineering, IEEE Transactions on*, 55(1), 41–50.
- Neve, M., Nicolao, G.D., and Marchesi, L. (2007). Nonparametric identification of population models via gaussian processes. *Automatica*, 43(7), 1134 – 1144.
- Nicolao, G.D., Sparacino, G., and Cobelli, C. (1997). Nonparametric input estimation in physiological systems: Problems, methods, and case studies. *Automatica*, 33(5), 851 – 870.
- Pillonetto, G., Chen, T., Chiuso, A., Nicolao, G.D., and Ljung, L. (2015). Regularized linear system identification using atomic, nuclear and kernel-based norms: the role of the stability constraint. *submitted to Automatica*. URL <http://arxiv.org/abs/1507.00564>.
- Pillonetto, G. and Chiuso, A. (2015a). Tuning complexity in kernel-based linear system identification: the robustness of the marginal likelihood estimator. *Automatica*, 58, 106 – 117.
- Pillonetto, G. and Chiuso, A. (2015b). Tuning complexity in regularized kernel-based regression and linear system identification: the robustness of the marginal likelihood estimator. *Automatica*, 58, 106 – 117.
- Pillonetto, G., Chiuso, A., and De Nicolao, G. (2011a). Prediction error identification of linear systems: a nonparametric Gaussian regression approach. *Automatica*, 47(2), 291–305.
- Pillonetto, G. and De Nicolao, G. (2010). A new kernel-based approach for linear system identification. *Automatica*, 46(1), 81–93.
- Pillonetto, G., Dinuzzo, F., Chen, T., Nicolao, G.D., and Ljung, L. (2014). Kernel methods in system identification, machine learning and function estimation: a survey. *Automatica*, 50(3), 657–682.
- Pillonetto, G., Quang, M.H., and Chiuso, A. (2011b). A new kernel-based approach for nonlinear system identification. *IEEE Transactions on Automatic Control*, 56(12), 2825–2840.
- Prando, G., Chiuso, A., and Pilonetto, G. (2014). Bayesian and regularization approaches to multivariable linear system identification: the role of rank penalties. In *Proc. of IEEE CDC*.
- Prando, G., Pilonetto, G., and Chiuso, A. (2015). The role of rank penalties in system identification. In *Proc. of SYSID, to appear*.
- Quenouille, M.H. (1947). A large-sample test for the goodness of fit of autoregressive schemes. *Journal of the Royal Statistical Society*, 110(2), pp. 123–129.

- Rasmussen, C. and Williams, C. (2006). *Gaussian Processes for Machine Learning*. The MIT Press.
- Riley, J.D. (1955). Solving systems of linear equations with a positive definite, symmetric, but possibly ill-conditioned matrix. 9(51), 96–101.
- Rissanen, J. (1978). Modelling by shortest data description. *Automatica*, 14, 465–471.
- Robbins, H. (1951). Asymptotically subminimax solutions of compound statistical decision problems. *Berkeley Symposium on Mathematical Statistics and Probability*, 131–149.
- Robbins, H. (1956). An empirical Bayes approach to statistics.
- Robbins, H. (1964). The empirical Bayes approach to statistical decision problems. *Ann. Math. Statist.*, 35(1), 1–20.
- Romeres, D., Prando, G., Pilonetto, G., and Chiuso, A. (2015). Classical vs. bayesian methods for linear system identification: point estimators and confidence sets. In *Submitted to IEEE CDC, available in ArXiv*.
- Rosasco, L., Tacchetti, A., and Villa, S. (2014). Regularization by early stopping for online learning algorithms. Technical report, arXiv:1405.0042v1.
- Saitoh, S. (1988). *Theory of Reproducing Kernels and its Applications*, volume 189 of *Pitman Research Notes in Mathematics Series*. Longman Scientific and Technical, Harlow.
- Sanandaji, B., Vincent, T., and Wakin, M. (2011). Exact topology identification of large-scale interconnected dynamical systems from compressive observations. In *ACC 2011*, 649–656.
- Schölkopf, B. and Smola, A.J. (2001). *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*. (Adaptive Computation and Machine Learning). MIT Press.
- Schwarz, G. (1978). Estimating the dimension of a model. *The annals of statistics*, 6(2), 461–464.
- Shah, P., Bhaskar, B., Tang, G., and Recht, B. (2012). Linear system identification via atomic norm regularization. In *Proc. of the 51st IEEE CDC*, 6265–6270.
- Shiller, R.J. (1973). A distributed lag estimator derived from smoothness priors. *Econometrica*, 41(4), pp. 775–788.
- Söderström, T. and Stoica, P. (1989). *System Identification*. Prentice-Hall.
- Songsiri, J. and Vandeberghe, L. (2010). Topology selection in graphical models of autoregressive processes. *Journal of Machine Learning Research*, 11, 2671–2705.
- Stein, C. (1956). Inadmissibility of the usual estimator for the mean of a multivariate normal distribution.
- Stein, C. (1981). Estimation of the mean of a multivariate normal distribution. *The Annals of Statistics*, 9(6), 1135–1151.
- Strawderman, W. (1978). Minimax adaptive generalized ridge regression estimators. *Journal of the American Statistical Association*, 73, 623–627.
- Sundin, M., Chatterjee, S., Jansson, M., and Rojas, C. (2014). Relevance singular vector machine for low-rank matrix sensing. In *Signal Processing and Communications (SPCOM), 2014 International Conference on*, 1–5.
- Suykens, J.A.K. (2001). Nonlinear modelling and support vector machines. In *Instrumentation and Measurement Technology Conference, 2001. IMTC 2001. Proceedings of the 18th IEEE*, volume 1, 287–294 vol.1.
- Sznaier, M. (2012). Compressive information extraction: A dynamical systems approach. In *Proc. of SYSID*, 1559–1568.
- Tempo, R. (1995). Worst-case optimality of smoothing algorithms for parametric system identification. *Automatica*, 31(5), 759 – 763.
- Tiao, G.C. and Zellner, A. (1964). Bayes’s theorem and the use of prior knowledge in regression analysis. *Biometrika*, 51(1/2), pp. 219–230.
- Tibshirani, R. (1994). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society, Series B*, 58, 267–288.
- Tikhonov, A. and Arsenin, V. (1977). *Solutions of Ill-Posed Problems*. Washington, D.C.: Winston/Wiley.
- Tipping, M. (2001). Sparse Bayesian learning and the relevance vector machine. *Journal of Machine Learning Research*, 1, 211–244.
- Verhaegen, M. and Hansson, A. (2015). N2sid: Nuclear norm subsystem identification. Technical report, arXiv:1501.04495.
- Wahba, G. (1990). *Spline models for observational data*. SIAM, Philadelphia.
- Wahlberg, B. (1991). System identification using laguerre models. *Automatic Control, IEEE Transactions on*, 36(5), 551–562.
- Wald, A. and Mann, H.B. (1943). On the statistical treatment of linear stochastic difference equations. *Econometrica*, (11), 173–220.
- Whittaker, E. (1922). On a new method of graduation. *Proceedings of the Edinburgh Mathematical Society*, 41, 63–75.
- Whittle, P. (1953). Estimation and information in stationary time series. *Arkiv för Matematik*, 2(5), 423–434.
- Wipf, D.P. (2012). Non-convex rank minimization via an empirical Bayesian approach. 914–923. AUAI Press.
- Wipf, D., Rao, B., and Nagarajan, S. (2011). Latent variable Bayesian models for promoting sparsity. *IEEE Transactions on Information Theory (to appear)*.
- Wold, H. (1949). A large-sample test for moving averages. *Journal of the Royal Statistical Society. Series B (Methodological)*, 11(2), pp. 297–305.
- Yang, Y. (2005). Can the strengths of AIC and BIC be shared? A conflict between model identification and regression estimation. *Biometrika*, 92(4), 937–950.
- Zadeh, L.A. (1956). On the identification problem. *Circuit Theory, IRE Transactions on*, 3(4), 277–281.
- Zorzi, M. and Chiuso, A. (2015). A Bayesian approach to sparse plus low rank network identification. Technical report, arXiv:1503.07340.
- Zorzi, M. and Sepulchre, R. (2014). AR identification of latent-variable graphical models. Technical report, arXiv:1405.0027v2.
- Zou, H. (2006). The adaptive Lasso and its oracle properties. *Journal of the American Statistical Association*, 101(476), 1418–1429.